LEVEL

AD A077991

Current Approaches To Decision Making In Complex Systems. II.

Report of a Second Conference held at Richmond, Surrey, England,
10-12 May 1976

by
G. Pask

SYSTEM RESEARCH, LTD.
37, Sheen Road
Richmond, Surrey, England

ARI LIAISON OFFICE, EUROPE
European Research Office, U.S. Army, London, England

MARCH 1978

Contract DAERO-76-G-043

Prepared for

DDC
RECEIVED
DEC 12 1979
D

DDC FILE COPY

79 22 5 100

U. S. ARMY RESEARCH INSTITUTE

FOR THE BEHAVIORAL AND SOCIAL SCIENCES

A Field Operating Agency under the Jurisdiction of the
Deputy Chief of Staff for Personnel

JOSEPH ZEIDNER
Acting Technical Director

W. C. MAUS
COL, GS
Commander

Research accomplished
by grant from the Department of the Army

System Research, Ltd.

CURRENT APPROACHES TO DECISION MAKING

IN COMPLEX SYSTEMS: II

G. Pask

System Research, Ltd.
37, Sheen Road
Richmond, Surrey, England

ARI Liaison Office, Europe
European Research Office, U.S. Army, London, England

U.S. ARMY RESEARCH INSTITUTE FOR THE BEHAVIORAL AND SOCIAL SCIENCES
5001 Eisenhower Avenue, Alexandria, Virginia 22333

Office, Deputy Chief of Staff for Personnel
Department of the Army

March 1978

Current Approaches to Decision Making in Complex Systems

Index

****

ABSTRACT

The proceedings of the Second (1976) Richmond Conference upon
'Current Approaches to Decision Making in Complex Systems'
held at Richmond are reported. For preference, the report
should be read in the context of an earlier report 'Current
Scientific Approaches to Decision Making in Complex Systems'
which sums up the proceedings of the First (1975) Conference
on the theme, issued as an ARI technical report in April 1976;
Contract DAERO-75-G072.

The following main conclusions were reached.

(a) Some individuals at least can be trained as
decision makers, but this should not be equated with training
in probability assessment, simple laboratory choice selection,
judgement, or the like.

(b) There is no uniquely "best" decision training
method but quite a lot can be said about techniques that do
not work on their own (though they may contribute to a training
process). The caveat of (a) is important. Decision is a
complex process, not a simple act as often supposed in laboratory
studies, so that training expedients like providing immediate
knowledge of results feedback, relative to a particular decision
(for example, objective probability) do not teach decision even
though they improve certain relevant behaviours. Effective
decision training involves at least remedying mistakes and
suggesting cogent heuristics in terms of a common language
established, perhaps, as a compromise, and understood by the
trainee and the instructor.

(c) Decision making is probably a collection of
skills integrated in the framework of tasks between which there
is adequate generalisation.

(d)   In specific situations decision skill can be
improved.   Probably there is sufficient task generalisation
and transfer of learning to comprehend classes of skills.

· (e)   Decision makers must <u>learn</u>, preferably in
context, concepts, facts, principles, attitudes, procedures
and heuristics (all the possibilities cited); some of them are
usefully <u>taught</u>.

(f)   "Learning to Learn" applies to a decision
situation (or class of situations).   "Learning to learn"
involves (amongst other things) formulating task representations
and problems.

(g)   The "optimum balance" of aiding against training
appears to be task dependent and dependent upon idiosyncracies
of decision style.

(h)   Several coherent structuring schemes were
discussed.   The environment may also be personally structured
by individuals who have 'learned to learn'.

The following issues were considered worth examining in the
future and regarded, by most participants, as inadequately
explored by current empirical studies and theoretical developments.

(i)   Prescriptions for training that involve problem
formulation and recognition in large but specific task settings.

(j)   Logics of action.

(k)   Structuring, both internal and external.

1

PREFACE

This report upon the Second (1976) Richmond Conference upon "Current
Approaches to Decision Making in Complex Systems" should be read in
the context of an earlier report "Current Scientific Approaches to
Decision Making in Complex Systems" which sums up the proceedings of
the First (1975) Conference on the theme, issued as an ARI technical
report TR-76-B1 in April 1976: Contract DAERO-75-G072. There are several
reasons why this is so; many of the participants attended both
conferences and continuity is established; the editorial process has,
so far as possible, removed the inevitable redundancy due to the fact
that old standing participants often spent effort in acquainting their
incoming colleagues with points of view that had already been rehearsed;
the background setting and style of the 1976 conference was very similar
to that of the 1975 conference (without this information it is impossible
to gain full benefit from the proceedings set out below). In addition,
it is desirable to read, also, the report of a conference in Oregon,
held shortly afterwards, entitled "The Training of Decision Makers"
by Drs Goodman, Fischoff, Lichtenstein and Slovic, since appreciable mutual-
ism exists; some participants attended both the Oregon and Richmond Conferen-
ces; the orientations were different but a great deal can be gained from
examining their more numerous common features, for example, the distinctions
between "abstract" decision theory and "task oriented" decision making (how
different are they if there is transfer of learning between decision
tasks) and the inseparability, in complex decision making and decision
analysis of "choice" per se, as compared to learning, conation, cognitive
operations and the behaviours in which they are manifest.

The organisation and administration of the Conference were due to
Elizabeth Pask, assisted by Linda Barsby. The discussion was tape
recorded and multiple tapes were prepared by A.I. Films Ltd.,
12A, Archway Street, London, SW13. Mrs Renee Bird, who transcribed the
proceedings, was present throughout the Conference to avoid ambiguities,
so far as possible, and the material was initially edited by Isaac Haissman.

2

INTRODUCTION

The setting of the Second Richmond Conference closely resembled
the setting of the First Richmond Conference. Both took place at
the Richmond Gate Hotel: position and technical papers mingled with
discussion. Whereas the First Conference was aimed at determining
the State of the Art (so that restrictions were minimised)
participants at the Second Conference were presented with several
specific questions.

(a)   Can individuals be trained to be effective decision-makers
in unstructured (and structured) situations?

(b)   If the answer to question (a) appears to be "yes", how can
that training be accomplished most effectively?

(c)   Does it make sense to think of decision-making as a skill,
or as a collection of skills?

(d)   Can decision-making skill or skills be developed in a
sufficiently general way that they can be applied in a variety of
specific contexts or a limited set of contexts?

(e)   What is it that the decision-maker needs to be taught?
Concepts? Facts? Principles? Attitudes? Procedures? Heuristics?

(f)   Can decision-making skills be developed through a "learning
to learn" type exposure.

(g)   Decision-aiding is a complementary approach to improving
decision-making. Insofar as an aid succeeds in simplifying or
otherwise facilitating the performance of some specific task, its
existence may lessen the training demands vis-a-vis that task.
On the other hand, users of decision-aids must be trained in their
use. What is the sensible balance?

(h)   A corollary approach to both training and aiding is job or
task design. How should decision-tasks be structured?

The Conference participants are listed below, together with their affiliations or institutions (those marked with an asterisk presented papers).

Dr Katrin Borcherding   *
Sonderforschungsbereich 24
Sozial-Und Wirtschaftspsychologische
Entscheidungsforschung
Der Universitat Mannheim (WH)
W. Germany

Dr Berndt Brehmer   *
Umea Universitet
Psykilogiska Institutionen
902 47 Umea
Sweden

Dr Donald Broadbent   *
Department of Experimental Psychology
University of Oxford
South Parks Road
Oxford

Dr O. Gelman   *
Merkaz Klita 52B
Mevasseret Zion
Jerusalem
Israel

Dr Robert Goldsmith   *
Department of Psychology
Lund University
Lund
Sweden

                                               *
Dr. E. M. Johnson
Organisations and Systems Research Laboratory
US Army Research Institute for the Behavioural
  and Social Sciences
1300 Wilson Boulevard
Arlington
Virginia 22209
USA

Dr Michael Kaplan
Chief ARI Behavioural and Social Sciences Liaison Office
European Research Office
223 Old Marylebone Road
London

Dr J.B. Kidd   *
Department of Industrial Administration
The University of Aston, Birmingham
Maple House
158 Corporation Street
Birmingham

Dr Gernot D. Kleiter   *
Psychologisches Institut
Universitat Salzburg
5020 Salsburg, Akademiestrasse 22
Austria

Dr A.G.S. Pask
System Research Ltd
Woodville House
37 Sheen Road
Richmond
Surrey

Dr L. Phillips   *
Department of Psychology
Brunel University
Kingston Lane
Uxbridge
Middlesex

Dr R.E. Schaefer   *
Sonderforschungbereich
Sozial – und Wirtschaftpsychologische
Entscheidungsforschung
Der Universitat Mannheim (WH)
Germany (BRD)

Dr C. Vlek   *
Institute of Experimental Psychology
University of Groningen
Biologisch Centrum
Kerklaan, Haren (Gr.)
The Netherlands

Dr J. Zeidner
Director
Organisations and Systems Research Laboratory
US Army Research Institute for the Behavioural
   and Social Sciences
1300 Wilson Boulevard
Arlington
Virginia 22209
USA

A full transcript of the proceedings, together with other
additional material (papers, position statements and so on) is
lodged at the Army Research Institute for the Behavioural and
Social Sciences in Washington.

Groups of presentations or topic series are prefaced by brief
overview statements: thereafter, the material is an abbreviated,
edited transcript.

Although the theme of training and learning permeates all of these
topic series it is convenient to classify them, albeit arbitrarily,
as follows :

1. Orientation.
2. Cognitive operations in decision taking.
3. Gaming and simulation, games as training systems.
4. Application of Specific Decision Theoretic Methods.
5. Logics of action for command and control systems.
6. Conclusions reached and general recommendations.

The report can be read sequentially since the length is considerable;
an alternative and probably more realistic reading scheme is to look
first at all of the "Overviews"

1. Orientation

Overview

The first (1975) Conference was given a direction by Dr Baker's
paper on the complex command and control simulation system SIMTOS
and to some extent, also, by Dr Stael Von Holstein's paper on
Decision Analysis as a discipline.

The present (1976) Conference was given a direction by the initial
questions and topics, by Dr Kaplan and by a context setting paper
from Dr Johnson, which spelled out the reasons why such questions
were asked. This paper provoked lively discussion which is reported
in some detail.

SECTION 1

Psychological and Technical Background to the Conference

Johnson: I want to provide some background with respect to the
general questions which have been posed, a brief review of what
we see as the state of the art with respect to the questions and
provide an overview of some of the directions our own research is
taking at the moment.

Although literature on decision making is extensive, there are
few studies which deal explicitly with the problem of training
and decision making skills. One challenge is to translate the
results of research in decision into its implications for training.

We need to distinguish between limitations of performance around
which the system must be designed and deficiencies in performance
which can be trained out of the system. One example of basic
limitations is the number of alternatives which a decision maker
can consider at any one time. One example of limitations which
can be trained out is the calibration of probability estimation.

First we need to consider whether individuals can, in fact, be
trained to be effective decision makers in unstructured as well
as structured situations. The answer to this question appears
to be "yes".

How can this training be accomplished most effectively? Does it
make sense to think of training for decision as a skill or as a
collection of skills? Can decision making skill or skills be
developed in a sufficiently general way so that they can be
applied in a variety of specific contexts. I think the question
of generalised ability is at the heart of training, in this case,
since its very clear from a variety of good evidence that we can,
in fact, train people to be effective decision makers in very
narrow contexts.

Other questions are whether the decision maker needs to be taught
concepts, facts, principles, attitudes, procedures, heuristics?
And can decision-making skills be developed through a "learning
to learn" type exposure.

Decision aiding is a complementary approach to improving decision
making. Although the focus of this conference is on training,
insofar as an aid simplifies or otherwise facilitates the
performance of some specific task, its existence may lessen the
training demands vis-a-vis that task. On the other hand, users of
decision aids must be trained in their use. There are several
examples of decision aids which have been in fact provided within
military settings which have never been used because the users
don't trust them or refused to use them: of "skeletons which have
been buried and no one talks about any more."

Learning is an invariable and unavoidable consequence of decision
making and probably of any other intellectual activity as well.
However, there is a feeling around many, particularly within the
defence community, that an individual who applies focus within the
fully trusted with crucial decisions until they have had some
experience with operations at lower levels. The operational
experience is often obtained from field exercises and manoeuvres.
A critical question concerns how much experience is required.
If the need for operational experience is deemed to be a require-
ment, for whatever reason, the immediate next question that is
posed has to do with implementation. Does this experience have
to be provided by command itself, or decision making in practice,
or can it be provided by a simulation.

Up to now the most frequent answer to this question in the defence
community has been one which relies most heavily on command
experience, perhaps because the tasks involved are critical to
success and the highest fidelity should be provided. But data are
not available to determine how much fidelity can be sacrificed.
Economic considerations, of course, enter the picture. It is
questionable whether the defence community can afford the luxury
of using field exercises and manoeuvres in providing this
operational command experience.

Morgover, in peacetime, an army has a very peculiar role. Commanders
and decision makers get graded on issues that are nothing to do
with the primary role of the military. They get graded on whether

the men are fed, whether the pay is on time and whether the ground are free. One illustration is that one of the most popular courses which is given at the US Army's Command and General Staff College is on unit- administration.

I'd like to provide a context for military decision making which I think is a little broader than is generally considered in behavioural science. In a sense, decision making is a psychological "waste basket". If you can't call an act of cognition or behaviour anything else, you call it "decision making". However, decision making is usually formulated as selection among alternatives or allocating the resources under decision maker's control, versus action, which is the commitment of resources and the implementation of decisions.

It is useful to separate decision making into three levels. First, strategic decision making which is policy definition. Next, tactical decision making, which is concerned with actual resource allocation, and finally, operational decision making, which is concerned with implementation. An example of this taxonomy which Ward Edwards provided, is in the context of mountain climbing. A strategic decision would be to decide which mountain to climb. A tactical decision would be to decide how you are going to climb that mountain. The operational decisions would be made by the individual who actually ends up climbing the mountain.

Each level of decision may be made by a different individual, or by the same individual. Military and industrial breakdowns are often different with a bias to distributing the load over different individuals.

Tactical decision making is the area of most concern for us and it encompasses a wide and diversified set of tasks. Military decisions have a number of characteristics which I think differentiate them from decision making in general. First, the criticality of the consequences, or the amount of risk involved. Next is the criticality of time as a dimension for effective decision making. A decision which is made too late is ineffective regardless of the quality of the decision itself. Military decisions have a dynamic nature.

There are relatively few static decisions in the military. In
contrast to the popular stereo-type the military organisation
often appears much more flexible in its decision making process
than does management or industrial decision.

First, the institutional framework within which military decisions
are made is highly structured. Major military decisions are normally
the result of staff actions where channels of information processing
are well defined. When I speak about training commanders I also
mean this as interchangeable with "command groups". Commanders
or decision makers in the military organisation rarely work in isolation.
They are supported by a command group or a staff with well
defined roles. For example, there is an intelligence officer who is
concerned with information gathering, an operations officer who is
concerned with different concepts of operation, a logistics officer
concerned with logistics, resources available, and so on. Although it
is quite highly structured, this arrangement corresponds, I think
very closely, to industrial decision making in that few industrial
managers make decisions by themselves. They are generally supported
by a staff, although in many cases the names and the terminology are
different. For example, an intelligence officer in a military
setting is analogous to a marketing officer in an industrial
company who obtains information on various markets.

Second, control of information throughout this command structure is
an accepted principle of operation. The decision making at lower
echelons is usually devoid of higher level information. Higher
echelons are restricted by the information which they can obtain
from lower levels. In other words, there are multiple lines of
information available. Higher echelons will receive information
which will never be made available at lower levels, as well as
having to rely on that information which is processed by the lower
levels.

Third, standard operating procedures and customs are well established.
Both established procedures and the interpretation of these procedures
by the commander and his staff provide a latitude within which
decision functions are exercised. That is, there are well structured

procedures for making decisions (which in many cases correspond very closely with those procedures which have come to be known as "decision analysis"), and these are the procedures which are generally followed. They contain criteria for determining the effectiveness of military decision making; but the evaluation functions are varied and intricate, and depend on a shifting frame of reference. For example, in peace time training, safety and cost may be the criteria, where in various forms of conflict other criteria become important.

Training is an area in which the military in general has become highly involved and I want to briefly go through the current 5-phase procedures used in the military for designing instructional systems. The first 3 phases, I think, are perhaps the most important.

The first phase is concerned with identifying what skills you want to train (the task performed in a particular job, those subtasks which you've selected for training) determining job performance measures for each task, and the selection of an instructional setting for the task. This phase is crucial in any consideration of training decision makers as it is not at all clear that we know what it is that decision makers do. If you wish to train decision makers what in fact are the performance objectives which you are trying to obtain? How is it that you know when you've trained the decision maker?

The second phase, "design", becomes even trickier in the sense that we need now to establish what our learning objectives are, determine how to measure achievement of each learning objective and some test of entry level behaviour. If we are not at all sure that we know what the learning objectives (of decision making skills in general) are, the terminal level behaviour will be inadequately specified and how can we know what the entry level behaviour should be.

The remaining three phases follow rather logically; development of instructional material, actual implementation, and then control which is the evaluation and revision of the instructional material. This is the general guide line, within which all the training material now used is being developed. In many cases these guide lines are more honoured in principle than in the actual implementation, since even a cursory glance shows that it is tremendous work to design and structure via the procedures as laid down.

What I'd like to turn to next is how one might go about training decision makers. Assume that we can break up the learning of a specific job or task into generic classes of learning addressed to job or task categories. We need to ensure that the student has the knowledge to identify the problem, generate reasonable solutions and evaluate these solutions. One implication is that you can't really teach the process of decision making without teaching it in some specific context. I don't think you can teach decision making in a vacuum, you teach decision making about something, which brings up the point of generalisability (positive transfer of expertise, from context to context).

Developing alternative solutions is perhaps the most difficult aspect. It is something we call problem structuring, getting students to structure problems into a framework within which they can use some of the concepts and tools of decision making. How we teach this is much like the "sorcerer's apprentice", and I think follows very much traditional lines, in that training in complex skills is traditionally being carried out by using a master/ apprentice relationship. In the absence of clear rules the skill must be passed on by example, and we're not exactly clear what it is we're teaching the students to do. Only that they do it the way we do it.   I think that graduate school is a good example of this. One of the questions is how does the learning take place. And sooner or later the beginner or the apprentice acquires some inner representation of the skill and learns to recognise it when it appears. He may still not be able to verbalise what it is that he has learned.

Last potential criteria we might use include : "predictability",
the problems viewed as if they were the same in reaching solutions' ;
"persistence", use of a favourite solution; "timeliness",
appropriately timed execution; "completeness" and "consistency".

The process of decision, if it can't be taught separately from
the content concerning which decisions are to be made, raises
some questions about how one should go about structuring the content
of the material to maximise decision making skills.

Another approach is "Decision Analysis". Several assertions can
be listed which pretty well summarise decision analysis. If you
believe these assertions (for example, decision making requires the
assessment of both values and probabilities), then you're probably
a good decision analyst and can probably teach students to believe
in these assertions. Students are generally rather gullible.
However, none of these assertions offer any clues for how you
might go about training people to actually employ them and I'm
not at all sure that this will make them better decision makers.

At this point we're not sure what gaming, which is something we call
game-base learning and the use of gaming as a decision making skill,
teaches the player. However, we are sure that we can get the
students to play, which is not at all bad as some of the other
forms of instruction have very low acceptance by students. Now
game based learning seems to have gone through three phases. The
first phase, acceptance on faith, widespread enthusiasm, and little
hard evidence. The second, post honeymoon phase, is disillusionment.
It was believed that games motivate, but no substantive evidence
exists that they induce critical thinking any more effectively than
other methods. The third phase is characterised by a real optimism
in which a number of games are evaluated in a wide variety of
settings. A pool of data on the learning effects of specific games
is accumulated and a clarification of the theory underlying claims
for what games can do as training media.

The reason I mention these 3 phases is that although concerned
only recently with the problem of training decision making using
games, we've gone through all 3 phases. At first we thought
games were effective and were very enthusiastic. After a little
more research on games we decided that games are a waste of time.
I think now we're in the optimistic phase; a number of games are
under development being tried out by a number of different players
in different situations. Two of the things we think are
being learned in games are the relationships between variables,
and the relationships inherent in the system. By being able
to observe the effects of past decisions, in a sense the decision
maker has to live with the results of a sequential series of
decisions unlike many of the static decision tasks. However, the
question of what is learned is difficult to answer because

different strategies may have different degrees of transferability.
For any given game, a number of different solution strategies may exist,
so that different strategies may place different allocations of
effort on different capabilities of the players and different
learning processes may be required for acquiring these different
strategies. The implication is that in any complex
learning environment, like games, an instructional
strategy is required to ensure transferability of the strategies
being learned by the players (in the sense that you don't want the
players to learn just that game, but hopefully some strategy which
has a degree of transferability). When I say "you don't want the
players to learn just that game", it is, in part, because we're not
clever enough to design games for which the only strategies are
those which we wish the players to learn.

Comparisons between games and other forms of training are
particularly difficult, though the difficulty is common to all
forms of training or at any rate to all forms that are suitable
for training decision makers and the root of the difficulty is
really the question "how do you measure the performance of the
decision maker"; and its one which we don't really have a good

handle on. We are not quite sure what criteria you use, to
distinguish good decision makers from bad decision makers,
irrespective of how weak the scale of measurement. A good
summary of current research and concepts for training decision
makers is a statement in a review by Dr Lichtenstein and others,
on the calibration of probability estimation in which they said

'The most striking aspect of literature reviewed here is its
dust bowl empiricism. Psychological theory is largely absent
either as motivation for the research, or as an explanation of
the results, as much of the research seems motivated by simple
questions concerning what would happen, so much of the interest
in research is in its potential applications.'

Vlek : I just want to express my support, for the remark that
decision making is more or less a psychological waste basket.
If you don't know what to call a piece of behaviour,
you call it decision making. On the one hand, that
speaks to my imagination. On the other hand, I think we could
also interpret this sentence in a more positive way because I
believe that the statement may be understood exactly the other
way round, namely that decision making is so pervasive in most
of human behaviour, that we haven't yet laid out all the elements,
all the separate human functions that are involved, and we haven't
obtained a clear idea of how people can be trained to exercise
these functions.

We might like to think of decision making as, maybe, an episode
in a larger process called problem solving, but that's not enough.
After we have made decisions, we have to implement actions. We
have to implement them according to plan, or we have to make
provisions for a modification on the way. Whilst implementing
actions we should be willing to learn from consequences that
occur and my idea is that we should try to lay out the whole
process of what happens both before and after the point where
someone makes a decision. That is not something to be thrown
into the waste basket".

<u>Johnson</u> : I was using "waste basket" in the sense of "being
pervasive", as well as pejoratively. Many phenomena which are
not well understood tend to be called decision making, not because
they're pervasive but because its a handy term for labelling events
or behaviours.

So far as the processes of decision making are concerned, many
taxonomies can be used for the kinds of skills involved in
decision making. Taxonomies of decision skills ranging from 6 to
35 different skills have been developed.

<u>Pask</u> : Decision making can be regarded as a collection of skills.
It isn't <u>necessarily</u> a waste bin (it may be under some circumstances),
but it <u>is</u> a hotch potch. Rather than there being one skill called
"decision", there are many and the question of training hinges upon
positive transfer of learning between them.

No doubt decision taxonomies have more or less merit in different
contexts, but I am dubious about the notion of classifying in any
absolute manner. We seem to be adopting the view that decision is
a facet of some larger process, (unlike the classical approach to
decision making which took the situation as being paradigmatic; for
example, choice amongst alternatives, alternative actions or
whatever, or even in its most liberal form as an aspect of
problem solving). As you put it a moment ago, rightly I think,
decision is more than that. It entails problem formulation,
learning, will or connation: the whole spectrum of cognition
and the life history almost of the individual, or the group or the
team concerned. If so, then surely there is a case for classifying
mental operations of which "decisions" are context (task) dependent
reflections, rather than classifying "decisions" as such...

<u>Vlek</u> : My feeling is that good decision makers should have an
appropriate model of the task environment in mind, and a basic
understanding of the underlying dynamic mechanisms of a situation.
A great deal depends upon whether you look at a laboratory
experiment where subjects have to know how to obtain a proper
response, or whether it is a complex management situation.
Decision making really consists of manipulating this model of
the task environment so that the situation changes in the way
you want it to change. If you don't know the basic mechanisms of
the task and you only look at the superficial symptoms, then you
don't know how to operate on that situation. If you have a basic
understanding you know which level to pull in order to change
the situation in a desired way.

<u>Pask</u> : Hence your dynamic model of the task interacts with the
decision maker's model of his own mental activity and <u>this</u> model
(the <u>task</u> model also, perhaps) is modified by learning.

There is a fairly regular progression in learning. At first the
decision taker can <u>describe</u> a situation (perhaps <u>his</u> view of it;
perhaps in a standard manner). If presented with alternatives that
span this description he can make reasonable probability estimates
for selecting a correct alternative, but at this stage no method
for problem solving exists (no algorithm) and probably the
mechanism for the task is unknown.

Later on a method is acquired which, if successful, must be
compatible with the task mechanism. At this stage it is easier
for the decision maker to give a method than to use it. In other
words, you have an analysis, your personal model, of how the task
works and you might be able to say, "well the right approach is",
verbalise it, or represent it in some other language, yet not
necessarily be able to <u>do</u> it particularly well. For example, I
can verbalise how to multiply matrices, and since I'm not good at
multiplying matrices (very few people are, they get the I's and
J's mixed round), I have to reconstruct this method on each
occasion.

Finally, some people reach the stage of overlearning. For this particular art, just a few of them, (bookies at a racecourse, for example, have overlearned matrix multiplication). A method, in fact, a host of methods coexist. They can <u>do</u> the task but <u>cannot</u> <u>say how</u> they do it.

I have a hunch that these stages in competence bear especially upon the <u>flexibility</u> attached to decision making in the environment with which we are concerned. We're not concerned, primarily, with laboratory environments (though they are indeed extremely relevant to analysing the processes), nor with the relatively static management environments in which <u>most</u> management games or economic games are played, but primarily with a command and control environment where things are in a state of flux. Flexibility is ability to adapt to a state of flux and also to see the wood for the trees.

<u>Kidd</u> : Gaming is a very peculiar area and I would perhaps rather pedantically like to suggest that it is the taking of decisions that's important. I'm not sure about the use of the word <u>making</u> decisions. I don't understand what the making of decisions means, except in, perhaps, the search for benevolent futures. Even we ourselves take decisions incessantly. To get up, to drink, to take business decisions and see the outcomes.

Decision analysis is important as an academic exercise to try and understand the nature of it all, whereas business people just get on with it day by day. More of an operational approach perhaps.

<u>Johnson</u> : I think decision analysis is used quite extensively, not so much in operational decisions, but in those relatively static decisions where the stakes are very large. For example, there are companies in the United States where any decision which commits the company to over a million dollars has to be supported by formal analysis of why thats a good decision. Xerox is a good example,

which supports an "inhouse" decision analysis group
which acts as internal consultant to managers at that level to help
them to prepare these analyses. A decision exercise. There are a
number of consultant companies who make rather a good living out of
providing consultant services in this area.

Pask : Apropos of that, perhaps the main function of decision
analysis is covert.

The techniques of analysis (modelling making, or of evaluating various
alternatives and so on, or setting up utility scales of an appropriate
kind) are overtly stated. But the main use of decision analysis
once this load of techniques had been mastered, appeared to be to
sense the dialogue between the analyst and the participants -
a group of people designing an airport, or somebody concerned with
forest control, etc. Decision analysts seem to say, the great
merit of decision analysis is not the thing you teach,
but the analysis of dialogue and finding out what meanings are
attached to the variables like G. N. P. or the
internal models maintained by whoever is responsible for taking a
decision ultimately. The analysis is almost an interview or
interrogation, whereby the situation is reformulated and somehow
liberty is given to reformulate this situation, on the part of
the decision makers. This blurs the distinction between gaming
and decision analysis, because gaming could well allow for non
verbal conversation and externalise cognitive acts in behavioural
terms, more objective terms, than looking at stretches of dialogue.

Kidd: Can we go back one stage before that. We've mentioned
decision making aids, perhaps like the multi-attribute methods
and so on where the analyst has to be trusted. We've talked about
gaming, but how do you train someone into believing that he's not
omnipotent, but really needs help sometimes? The realisation or
perceptions, that he can't solve everything. Its simple. If you
have a large decision, you don't have to think about it, you go
out for help. But in many situations, perhaps, the decision maker
may say, "it is my job to take these decisions on my own", and,
"I'm somehow deflating my position if I go out for help".

Pask: When you talk about "needing aid", are you talking about the situation in which somebody needs aid to overcome these frailties and limitations? Or are you talking about the fact that a problem is so enormous that more information is needed? If that were the case, standard procedures exist for getting it, examining the data base in greater detail.

Kidd : I don't want to determine what aid is, just that a person has to recognise that he does need some aid. And you made an awfully big assumption at one stage by saying it is relatively easy to give this person training in sampling, design, searching data bases and so on. Whereas you accepted it seemed to be much more difficult to provide multi-attribute aids or Bayesian inference aids and so on because of their known technical problems as far as we are concerned as decision analysts.

The greater problems is further back at the first step.

Pask : You mean the framework of the research or the structural framework is inadequate.

Phillips : I think this is very much a problem with training. I do some lecturing to people who are remarkably unsympathetic to the quantitative aspect of decision analysis and have, as a result over two or three years, shifted the emphasis in teaching to try to show the quantitative road in order to get a better qualitative feel for the problem and I think particularly in the last year or so decision analysts are using sensitivity analysis more and more to show that this is the point where the quantitative influence in numbers, utilities, probabilities, importance weights, what have you, can be fuzzed out to see their effects on the decision itself.

Johnson: The emphasis on sensitivity analysis is, I think, one that's become very prominent since its very clear how shaky most of the numbers are that go into a decision analysis, irrespective of rigour. Point estimates in most cases really fail to capture all of the information. Sensitivity analysis turned out to be very useful in identifying those numbers which can, in fact, have very little influence on the problem or the alternatives and those numbers for which, perhaps, you need to re-examine the matter.

Phillips : Well, I'm also astonished at the number of people who mix up probabilities and utilities, who find this a new concept, who when you talk about the analysis of multi-attribute outcomes, are quite interested and amazed at the distinction between the importance of an attribute and the scale of entity, the location of an outcome on some utility scale. You really can talk about a lot of these things in purely verbal terms without getting too involved in the model and yet, when you talk to them, they say, "oh, you've clarified the decision making process, I understand it from a new point of view and a different point of view", and you can rightly say 'alright, in terms of any behaviour measured does it make any difference, in the long term, to their decision making". I can't answer that question.

There is a problem.   Can you imagine, here are people going on taking decisions for 20 or 30 years of their life, and then suddenly they get exposed to decision analysis for 3 days and then you expect them to take decisions differently.

Vlek : Apart from that, any decision analysis is meant to transcend cognitive limitations in habitual thinking about forming decision problems.

22

Phillips : I would hope that effective training would make it difficult to be aware of using decision analysis. I assume that decision analysis is a model that can be applied in any level of detail that you would like.

Vlek : Its equally important to know descriptive aspects of human functioning, because very frequently you don't have time to do an analysis, and then its extremely important to know what people tend to do naturally when they don't have time. What do you do? You concentrate on the most important utility aspect. For example, or you don't anticipate consequences farther away than 3 months from now. This is as important as being able to apply decision analysis procedures.

Decision analysis, like statistics, is just a method and it can only be used to its full potential when you have the time and other resources to hand.

2. COGNITIVE OPERATIONS INVOLVED IN DECISION TAKING

Overview

Four papers have been classified under this heading : those of
Drs Brehmer, Goldsmith and Vlek were presented and discussed in
sequence. Although Dr Phillips' paper was presented and discussed at
a later point (after the consideration of management gaming), it
properly belongs to this category and is thus included in the outline.

Between them, the papers exhibit a gradation of approach to the
cognition that underlies decision, along a dimension of macroscopic-
microscopic examination and a dimension of experimental method:
standard - specially devised.

Dr Phillips uses standard experimental methods for data gathering:
questions and inventories, for the most part. His model is also the
most macroscopic or "molar" and the research is largely directed to
determining intercultural differences in decision habits. In this
work, there is an assumption that individual modes of "probabilistic
thinking" or "dealing with internal states of uncertainty" are coloured
by the individual's culture and the typical thinking habits of a culture
(as revealed by questionnaires, etc) reflect individual habits: cultural
characteristics are regarded as parodies of or vignettes representing
the characteristics of individuals and the intercultural comparisons
may, in this sense, be regarded as comparisons between extremes of
individual decision style.

The model itself was presented at the 1975 Richmond Conference
and is written up in the 1976 report. The paper at the present Conference
reports further work along similar lines and additional studies bear
upon quite distinct fields, namely the personality profiles of
probabilistic thinkers and training in the assessment of probabilities.

(i)   Interesting differences and similarities of habit are
described and may be explained by recourse to the model.

SECTION 2

(ii)   That there is no correlation between personality variables and variables that satisfactorily account for probabilistic thinking (also true of some style-indicating variables).

(iii)   That verbal and numerical expressions of probability assessment are substantially independent.

(iv)   That there is a remarkable lack of positive transfer of training in the skill of probability-assessment, between distinct tasks.

(v)   That probability   assessment-training has anomolous features.

Whilst the model and its empirical support provide answers to some of the original questions (a to h in the Introduction), the model is insufficiently detailed to answer all of them. However, the author addresses himself, in general commentary, to all these questions.

Dr Brehmer's paper is set against the theoretical background of Brunswick's probabilistic functionalism and the "Lens" model for judgement. For this reason any particular interpretation is task specific ("cues", for example, depend upon a task model and a model for how decisions are made). With this qualification the underlying model is more microscopic or "molecular" than Phillips' and very definite experimental hypotheses can be tested on the assumption that subjects attend to certain cues and attempt, so far as possible, to obey certain rules. Moreover, the "Lens" model equations contain terms which refer to what a decision maker knows and other terms which refer to how he uses this knowledge to achieve a criterial performance; the paper is concerned, specifically, with the question of how knowledge is used and how people may be trained to use it better. At this level of description, several of the original questions can be answered for particular experimental situations, but because of the inbuilt task-dependency of the "Lens model" construct, difficulties arise over the extent to which decision habits, observed in the context of laboratory tasks, will govern behaviour in general. The problem becomes fairly acute when the laboratory tasks are standard (stimulus, response, rule for responding), whereas the tasks of interest

are an integral part of a complex and usually dynamic system.

Dr Goldsmith's paper describes a methodology and an associated model for cognition. The level of description to which the model is addressed could, depending upon the experimental design, be just as macroscopic or microscopic as the "Lens model". However, because the model is embedded in a participant-experimenter (alias the decision analyst) methodology, the situations are not standard, the description is a different kind of description, not committed to (though not excluding) the "Lens" formulation. The method involves two processes.

(A) Presenting a subject with a very complex decision task, and requiring the subject either to anticipate or immediately recall his strategies in reaching decisions.

(B) Refining this data to achieve a model, by means of an interactive decision analysis designed to eliminate misconceptions and to encourage the adoption of decision norms.

The result is a compromise which can be cogently modelled. The compromise model contains elements that stem from the subjects strategics and other elements due to the strategies recommended by the participant decision analyst.

Repetition or iteration of (A) (Process Tracing), and (B) (Analysis and Recommendation of Norms) yields, in convergent cases, a more refined compromise model specific to a partly subject determined task (ie. how the subject sees and comes to see the complex decision situation). Further, systematic iteration of Process (A) and (B) is proposed as a candidate training method in its own right.

The discussion of this paper contains material bearing upon jurisprudence, since this is one field in which complex and realistic decision situations are very commonly encountered and resolved. At first sight, the issue of "correct" or (distinctly) of "right/fair/ sensible" decision looks like a digression. In fact (and to some

extent in retrospect) debate on the topic of "truth" in contrast to "agreement" is germane to any complex decision process. It links closely with the "logic of action for command control systems" (Section 5 of this report) and necessarily comes into the foreground in the process of "compromise modelling" decision by means of steps (A) and (B) repeated systematically, either with respect of one decision task or many of them.

Dr Vlek's paper gives by far the most detailed account of decision. It lies at the "microscopic" end of the macro/micro or molar/molecular dimension and furnishes specific answers to each of the original questions (a to h in the Introduction). Decision is seen as involving all aspects of conation, cognition and behaviour (by way of contrast, the other participants, to a greater or lesser extent, conceived of decision as a special "faculty" interacting with mental processes in general). The underlying model is a "quasi hierarchical" structure for decision construed in this global manner, together with a taxonomy of variables that influence decision.

Dr Vlek's argument is complex and he prepared a revised and condensed paper which is included in this Section.

Several fairly definite points came out of the discussion over and above the (already noted) points (i) to (v) and these are briefly summarised below.

(vi) Decision is task specific . No decision takes place in a vacuum. The decision process involves the task in hand as well as the decision maker. However, the task situation is only given in part by the environment. Part of it is due to the decision maker's conception of the environment (this may be exteriorised, for example, the compromise model noted by Goldsmith).

(vii) Generalisation takes place in connection with positive transfer of training between complex and task oriented skills, familiarity with which is a prerequisite for decision though not decision itself.

(viii) The majority of experiments use laboratory tasks that are oversimplified (selection, confidence estimation) or only fully comprehended by able probabilistic thinkers (for example, lotteries) A great deal of confusion has arisen due to results derived from laboratory paradigms.

(ix) It was generally agreed that decision is a much more complicated process than most experimenters have imagined.

(x) Decision may, but need not, involve probabilistic thinking. Quite different heuristics are admissible as practically effective.

(xi) Decision should not be confused with response modes (selection, confidence estimation) . It may be a learnable skill or a cluster of skills between which positive transfer of learning can be achieved. It may be inseparable (Vlek) from the rest of

mental activity and behaviour.

(xii) Subjects can become facile in the use of response modes after suitable training and training can often be given by simple expedients (for example, providing knowledge of results feedback). In contrast, cogent decision training calls for a more positive and interactive approach (for example, remedying mistakes and proposing specific heuristics).

Whether or not such training is effective remains a matter for debate. When it is effective, the mechanisms involved are still quite obscure as there is little agreement upon the issue of whether general problem solving or decision is a uniform process, whether it involves clusters of skills, or whether decision is inseparable from the flux of conation, cognition and behaviour.

"Three Studies in Probabilistic Thinking"

<u>Phillips</u> : I'm not entirely sure what I mean by probabilistic
thinking, it is part of the research goal to get clearer about
it, but I do have three preliminary notions to begin with.

First of all, I mean by probabilistic thinking that a person
adopts a set to view the world as containing shades of grey with
regard to the occurrence of events. Some people say, "yes, this
is going to happen, or no, that's not going to happen". Others
say, "well there's some possibility it might happen". This kind
of person views the world in terms of uncertainty, when appropriate.

A second aspect of probabilistic thinking, is an ability to make
fine discriminations of uncertainty. Some people are better at
looking at that grey area between black and white, and some are
able to cut it more finely than others.

The third idea has to do with the ability to make meaningful
probability assessments. That is to say, to act as a measuring
instrument for one's own state of uncertainty and provide numbers
that have some bearing on the internal state of uncertainty.

Many more things could be meant by "probabilistic thinking" in
addition to those three. I would, for example, assume that the
last aspect of probabilistic thinking is where the heuristics that
Tversky and Kahneman talk about will operate. These heuristics are
personalised methods for taking internal states and translating
them, eventually, into communicable or numerical statements.

Let me describe three studies that we've conducted at the
Decision Analysis Unit at Brunel. The first one, people have
already heard about who were here last year; let me just give you
an indication of the sorts of things that we're doing. First of
all, I was very struck many, many years ago that there seemed to
be individual differences that I thought were related to culture,
having to do with an individual's ability to assess probability.
So we attempted to put together some instruments that might tap
these differences, and then set out to measure a probabilistic

thinking in different cultures. Let me tell you first about the
instruments. There are three kinds of instruments that we use.

(a)   The "view of uncertainty" questionnaire.

(b)   A card sorting task, and

(c)   Probability assessment questionnaire.

Now from these three tasks we derive about a dozen measures.

The view of uncertainty questionnaire is easy to administer and
consists of about 45 items like Fig.2.1. Subjects are asked to
write in the space provided a reasonable and appropriate response
to each question. The results are coded, since its a free response,
into a number of categories; the number of "yes" or "no" responses;
the number of "don't know" responses; the number of probability
words or phrases used; the number of different probability words or
phrases used; the "hit rate" or calculation of veridicality of
response for those questions that refer to facts and historically
known events; finally, a "catch all" category for things like,
"I hope so", which is, in our coding scheme, a statement of utility,
not a probability. The view of uncertainty questionnaire is meant
to tap a person's set to respond to a world in an uncertain way.

The second instrument is an attempt to get at this discrimination
of uncertainty. We have 29 cards on each of which is written a
word or a phrase, "like", "not likely", or "highly likely", or
"fairly sure". Subjects are instructed to sort the cards into
piles, so that all the cards in a given pile mean the same thing.
And the idea would be that a person who is not able to discriminate
uncertainty finally, would create a very small number of piles.
The measure that I use here is simply the uncertainty measure. H
(Fig.2.2). The bigger H, roughly the greater the number of piles.

The last task consists of a number of questions. Subjects tick
responses A or B, and gives an assessment of a probability between
50 and 100 to indicate how sure he is that he ticked the right
answer. If you haven't any idea, write 50, if you're dead sure,
write 100. By hypothesis this taps the uncertainty the individual
experiences in searching memory or working things out or going
through some kind of cognitive process to arrive at an answer to
these questions.

Calibration curves relate the probability that the individual assesses, which in this case would be between .5 and 1 and the overall percent correct, between 50 and 100 percent. A well calibrated probability assessor is somebody who stays along that 45 degree line; for example, somebody who, on all occasions when they say .8, turn out to be right 80% of the time. A couple of years ago I invented a model which manages to capture calibration curves in two parameters. See (Fig.2.3). The model and the report of the 1975 conference relates assessed probability to unbiassed probability. The family of curves are all generated by the logarithmic equation shown.

There is a problem about assessing the very best values of A and B, and we have a Bayesian inference program that does this. It finds posterior distribution across those two parameters, B, and, A, and then it searches for the mode of that posterior distribution.

One study is cross-cultural, involving both the basic measures and the calibration curves. I reported some of this last year but these are rather more detailed analyses. In Fig 2.2. I've indicated, as 1, 2, 3, 4, 5, the different cultural groups we've studied so far. On average the English students use about 25 yeses or nos, to the Chinese students 30-36. Chinese businessmen are fairly similar to the English students, Now I had hypothesised that the Chinese in fact would be less likely to adopt a probabilistic set than the English.

Very briefly, the reason I believed this is because I think that the tendency to adopt a probabilistic mode of thinking is very much tied up with the model of causality implicit in a culture. We are very used to a kind of Laplacian model, in which the world is deterministic, though not knowable completely. Everything has cause and effects, there are sequences of causes and effects. I can't know everything about these sequences and, therefore, I'm always in a state of partial ignorance and this state of ignorance can easily find expression as an uncertainty. The Chinese, in studying some of the tenets of Taoism and Confucianism and so

forth, have a different view, namely, that there is an ordering
principle, a grand principle that has created a sequence that's
in fact predetermined, a sequence of events, and these events flow
along under the order of this principle. Its not that one event
down here causes the next event, its an over-riding principle that
has "put things in their place and causes things to happen", if you
even wish to use the word "cause". I would argue then that a person
growing up in that culture would entertain the idea that you either
know what is going to happen next, or you don't, and if you know
what's going to happen next and your brother doesn't, then obviously
you can take advantage of this and perhaps that accounts for the
great keeness of gambling amongst the Chinese.

At any rate, the data are quite clear (Fig.2.2). We do get
differences and must account for them somehow, whether you like the
cultural hypothesis or not.

We purposely set out to find subjects who had some considerable
experience with English so that we're not testing people who have
only had brief acquaintance with it, though I suspect that this
just reflects a lesser variety of .... Now here's an interesting
one, it didn't get any differences in the H measure, that is to say,
the English and the Chinese subjects tend to sort the cards in about
the same manner. Not very much difference there. I suggest here
that although the Chinese are not likely to adopt a probabilistic
set in viewing the world that once you give them that set which
you've done after you've given them these cards because they have
all those little uncertain words written on them, they're quite
capable of making the discrimination. At least they're capable
of making verbally expressed discriminations. The reason I say
that is because, although I haven't shown it here, if you look at
the H measure for the probability assessment questionnaire, you
there discover the English students using a greater variety of
numbers than the Chinese. Many cross cultural investigators have
noted that numbers are a relatively culture free response mode,
and so I don't think that the difference between the Chinese and
the English is a matter of familiarity with the numbers, and the
fact that the English students use more, a greater variety of

numbers, than the Chinese students, therefore, giving a bigger H measure suggesting the rather odd conclusion that although the verbally expressed discriminations don't differ, the numerically expressed discriminations do, and that in turn suggests that the information processing that's going on is, in fact, a function of the response mode itself.

What I didn't show last year was the calibration curves. The calibration curve for the Brunel (English) students is the black line, Fig 2.4. The blue line is the calibration curve for the Howard and Peterson data that appeared a couple of years earlier. Now you can see that the calibration curves are at least roughly similar. The calibration curve for the Chinese students (Fig.2.5) is distinct. Doesn't matter what they say, they said, they are right about 50% of the time, whatever they say. Except when they say, 100%, and then they're right about 65% of the time, contrasted to this 75% of the time. These are quite clearly differences.

I also have done this on a group of University graduates whose identity I have to protect or the people who let me do this experiment would be very angry. They are highly selected University graduates who have been out of University for about two years; they are about 60% Oxbridge graduates, that is, they came from Oxford or Cambridge. Nearly every one of them has a non-science degree, history, government, politics, and so forth. So they are a very bright group, very capable, and their calibration curve for that same questionnaire is shown in Fig.2.4. Finally, in Fig 2.4, there is a calibration curve for a group of business men, for the most part working in various parts of Great Britain at branch offices of a very large company. Well, the clever thing about this is that we're using it in the following way.

We are also using the Joreskog approach to factor analysis which ensures proper solutions. All the other versions that you probably know can give some extraordinary improper solutions and some very wrong results in fact, and I'm using it on the English samples in an exploratory way to set up hypotheses which can then be tested

on the other samples. Now, first of all, let me point out that
I've run factor analyses on those 12 variables of the English
students, of the Chinese students, and for those subjects that
were the English graduates. Now the interesting thing is that a
six factor comes out for each of these; six factors are required
before you get a good fit and that is on only 12 variables.'. But
that is alright because, after all, I was purposely setting out to
find variables which I hope didn't have too much redundancy, in
the belief that they really indexed different aspects of
probabilistic thinking. But in each case we get six factors and
the other thing that pleases me greatly is that the factors are
almost identical for each of these different groups. The overall
results (still incomplete) are shown in Fig 2.6, (the factors come
out in slightly different order for each group, but, nevertheless,
they are roughly the same and I'm going to try to make some guesses
about that these factors might represent and you might help me).

Factor 1 loads highly on the number of one hundreds, the number of
correct one hundreds and the uncertainty measure on the probability
assessment questionnaire; these all turn out to be on one factor.
Notice it is a factor consisting entirely of numerical variables
and I suggest it has something to do with numerically expressed
discrimination of uncertainty.

Factor 2 has to do with the yes, no's, and don't know's and that's
purely verbal. It is verbally expressed things, written things,
not numerical. Don't know just stands out all by itself for some
curious reason. It has some loading on the next two factors, but
"don't know" seems to be a factor all by itself – the tendency to
say don't know.

Factor 3 is all verbal. Yes/no's, don't know's, number of probability
words and the number of different probability words. I would say
that this one does measure what I originally called "the probabilistic
set", but in addition to that it's also the verbal expression of
uncertainty. This is clearly a verbal probabilistic set and then
expression of uncertainty factor.

Factor 4 loads on hit rate 100, that one comes up in all of the
analyses except for the Chinese students. This is just totally
absent, this factor for the Chinese students. I don't know why
until I've done the hypothesis testing. But I call it "knowledge".

Factor 5 is both numerical and verbal and this one appears to be a
numerical expression of "don't know". It's the "how do you say
don't know in numbers". Because, you see, the number 50 loads
positively and the more frequently that you say 50, the less
frequently spread out your responses.

Finally, factor 6 is a calibration component. A and B never load
on anything else. A and B always come out as separate factors.

Interpretation of factors is always difficult and suspect. My
main theme is that factor analysis has, at any rate, shown that
life is more complicated than I thought it was. Probabilistic
thinking is just not a matter of two or three things; its a
matter of quite a few things. I don't too much care which factors
come out, but I am very interested to know that the factors that do
come out are fairly stable across groups and things that I thought
of to go together don't, and they don't for any of these groups,
and we've now got that conclusion on several hundred subjects.

NOTE : There was a great deal of discussion about the meaning
of the factors and Dr Phillips summed up.

Phillips : I feel as though 50/50 and 100 have a special status.
Psychologically they are not probabilities. The other thing that
comes out clearly is that these factors never, without exception,
mix numerical and verbal components. These are different. I am
suggesting that the response required (numerical or verbal) has a
significant effect on the cognitive processes each individual
employs in dealing with an internal state of uncertainty.

The next study involved a factor analytic comparison of variables
related to personality and cognitive style such as authoritarianism,
conservatism, tolerance of ambiguity, internal and external construct
and the 12 variables indexing probabilistic thinking.

When you do the factor analysis on the probabilistic thinking
and the personality variables,on the whole thing together, first
of all, the six factors already described for the other groups
emerge from these 12 same variables, which was greatly reassuring
because they were such a different subject population. Next, the
6 factors that emerge from my measures of probabilistic thinking
have no loadings with any of these variables. And the factors
that emerge from the measures of personality and cognition have
no loadings on any of the probabilistic thinking variables.
None, I really mean nothing.

Vlek : Do you have an explanation?

Phillips: Actually there is a kind of explanation. The person-
ality variables all have to do with utilities, nothing to
do with probabilities. These have to do with people's beliefs
and value systems and so forth, they don't have anything to do
with how people assign probabilities. They have to do with how
people assign values.

Broadbent : The internal/external control people would say it
was your assessment or dislike of uncertainty about things that
you are doing and the effect on the world, rather than about arbitrary
questions.

Phillips : Well, yes. That would be true of the overall scale,
but the subscales on internal/external control should have
loaded. For example, there's one thats just called the luck/fate
scale. There's another one has to do with political events.
Another one has to do with academic success. None of those have
any bearing at all on any of these probabilistic measures.

There are occasional little loadings;   maybe 0.3 or something
like that. But nothing beyond .4 at all. I was sure that something
would come out of it. There would be some correlations.

Pask : There is a more general probability question of distinguishing the notion of doubt, of entertaining doubt, and the notion of thinking probabilistically. It is interesting but it isn't altogether surprising that these personality and style measures don't have much relation to probabilistic thinking, as such. It is very surprising that you did not find some relation to doubt because I can't imagine anyone who doesn't entertain doubt. To be aware is presumably to entertain doubt.

Phillips : It seems very odd to me. At some level of cognitive mechanism there should be stability insofar as it does not matter whether you express ideas as a word or a number, but clearly it does seem to have some effect.

Vlek : Could I make another comment? I'm surprised about the unrelatedness between two sets of factors because I feel that with the dogmatism, conservatism scales you're tapping attitudes and your own explanation of the six factors in the probabilistic thinking study suggests that your six factors there, measure, use of response devices rather than attitudes and ways of probabilistic thinking in different ways which are inter-correlated so as to produce six factors. You measure a totally different complex of things.

Phillips : The purpose of the last experimental study was to find out whether or not you could train people to be better calibrated probability assessors and then whether or not the training would generalise to a different set of events. Each subject experienced and responded to four sets of events. In the first set of events the subject got an almanack questionnaire. He just went through all those 75 items. He didn't get any information at all about his performance. That was to establish a base line. The second questionnaire which is supposed to be roughly parallel to the first questionnaire, this time he received scoring feedback and some verbal training. That is to say my research assistant told him how he was doing. 75 questions again, score after every single question, and knowledge of results after each question as to whether A or B was the correct alternative; all of it was all computer controlled so that the cumulative, the score on that particular trial was displayed as

well as the cumulative running score and subjects were told
they would be paid extra according to cumulative score. Finally,
each subject received exactly the same questionnaire, again
with no feedback, just to see if there was effective training.

Each subject engages in a task series of the same design at
another session in which he gets no feedback. But this time
the task is to discriminate between a square and a circle
displayed on the screen of the CRT by the computer. You have
to tell me which one is larger in area. And then tell me how
sure you are. The one you've selected as larger in area.
These are all the almanack type questions of a binary nature
and this is a binary task as well, but it's a totally different
kind of task. The design is balanced for comparison of training
and generalisation of training over the types of task (Fig 2.7)
and some results are shown in Fig 2.8. With the geometric task
there is an appreciable training effect but not the almanack
task. But after subjects have had geometric training and are
asked to do the almanack questionnaire, they are no better off
than subjects who just started fresh with the almanack type of
questionnaire. No generalisation at all.

NOTE: Dr Phillips made points which he regarded as "the state
of the art", at various points in the discussion. These replies
are collected below.

Phillips : State of the art, yes, there seem to me to be quite
clear cultural influences on probabilistic thinking. What are
some of the unsolved problems, what on earth is the mechanism
that causes these. What is it about these cultures that causes
these differences, I don't know for sure. Future research seems
to me ought to be oriented toward finding out how these differences
come about. Insofar as we understand these differences and the
stability or non-stability of the factor structure from one

culture to the next, we ought to get a better idea of the cognitive
processes involved and that should apply within a culture as well.
So I've got increasingly keen on using cross cultural work to
highlight the nature of cognitive processes, the cognitive mechanisms
that are going on and hopefully are applicable within a culture.

As far as the effects of training in generalisation, I am convinced that
scoring rule feedback is the wrong way to approach it. The
reason I'm convinced its the wrong way to approach it, is that
because its like, its too crude, it doesn't really give any
information. Its too much in a kind of behaviourist tradition,
just letting the subject know where he's at. What you really want
to do is get into his head and say, well now, what kinds of
heuristics you are using, how you're doing this, are you doing
that. Much, much more complicated.

I don't know that my conclusions apply to anything
but scoring rules.   I think scored rule feedback is
just generally much too crude for that to be very effective.
And in fact, if you do look at any scoring rule research, you
rarely get very appreciable effects of scoring rule.

As far as the future is concerned, it does seem to me that its
rather necessary if we're going to talk about training, particularly
in probabilistic thinking to know what it is that we're talking
about. Is there any point in linking this to other cognitive
theory? It hasn't certainly been done very adequately so far,
and as far as relating it to measures of personality, that seems
to be a dead loss, yet it must relate to something, some other
aspects of cognitive thinking, unless we've identified some new
aspect of cognitive thinking that people haven't looked at
explicitly before, but I find that hard to believe.

The trouble with all this (regarding immediate corrective feed-
back) of course, is this kind of feedback is tantamount to saying
to the person, change your behaviour, but you don't tell them
how to do it.

In order to tell them how to do it, you've got first
of all to have some idea what he's in fact doing. So that's why I
have tried through the various studies we've been talking about to
get a clearer idea of how people order uncertainty in their head and
how they partition it off. I would expect, for example, that one of
the things that you've got to do in training is to get people to
adopt a probabilistic set when its appropriate. Because some people

just don't. I found very fascinating a study, I think it was by
Egan, which showed that when you change the difficulty of the single detection task, the person's calibration deteriorated but that was only one
person. Another person was able to adapt, to give probability
judgements in the two tasks that yielded exactly the same calibration
curve. Now nobody has attempted, as far as I know yet, to train
people to adopt some kind of fixed internal criterion. You see,
what's apparently happening is that subjects are shifting the
distribution of responses inappropriately. That is to say, they
figure perhaps that they have to distribute their responses between
50 and 100. And by keeping the distribution roughly the same between
an easy task and a difficult task, you get two different calibration
curves. What in fact you ought to do when confronted with a difficult
task is to shift the distribution of responses down the scale nearer
50/50. That's what I would train people to do. I'd try to get them
to be responsive to the difficulty of the task.

Pask: There seem to be three sorts of finding that emerge
from this. The first of these is I think Jim Baker did some work on
essentially task training. He was using a confidence estimation
technique with an automatic normalisation of scores and histogram
display, and what you've got here was the training in the task, and
this is observed, I think, if anybody has used that technique.
Dirkzwager observes the same training. You become able to handle
this situation and there is a training in being able to handle that
estimation situation, because its there in front of you.

This is not training people to calibrate better, it is just training
them to use the equipment or the response mechanism provided.

Another finding appears to be that it is useless to provide immediate
feedback if you are trying to 'recalibrate' subjects or to improve

their assessment. This appears to be one of your results, though your result on the geometrical task is, however, compatible with training to use a response equipment equally well.

Finally, there is, if you say how to improve or get into the head, then you can alter the calibration of some people (not all, perhaps), because, in fact, you're pointing out critical features either of a generalisable estimation skill or of the stimulus situation of a task situation, or both, by giving an explication of it.

Vlek : Theorise a little bit about what happens when you ask people for a probability statement.

Phillips : I think that the criticism to be levied against me is that I started with much too simple a conception of what was going on. I figured that uncertainty registered as some kind of feeling, and it was just a matter of my serving as a measuring instrument for that feeling and turning it into a number. This is kind of psychometric model, almost. It's dead wrong, I'm sure. I'm now convinced that the cognitive processes involved, in uttering this number, are vastly more complicated than I ever suspected.

Vlek : You end up with the feeling once you have gone through the process and then you want to translate it into a number because the experimenter requested that.

Phillips : I'm not sure, it may be. I'm amused by many of the memory theorists who have these nice block diagrams of all the things that are going on, retrieving things from memory and so on; and somewhere there is usually a little box that says "decision processes". Now I think that box is probably at least as complex as all the stuff that's outside.

A group of kangaroos is called
a. a troop
b. a pack

What is the capital of New York State?
a. New York City
b. Albany

Which ocean is larger (in sq. miles)?
a. Atlantic
b. Indian.

Fig 2.1.

		English students N = 42	Hong Kong Chinese students N = 53	Hong Kong Chinese Nurses N = 20	Chinese nurses in London N = 23	Hong Kong Chinese Business men N = 33
No. of Yes/No	M	24.8	30.9**	36.4**	30.1**	26.6
	S	9.3	9.9	9.1	8.1	10.1
	t		3.1	4.7	2.4	0.8
No. of Don't know	M	5.8	6.2	7.8††	8.6*	8.6*
	S	4.7	5.6	3.8	5.3	6.6
	t		0.4	-2.7	2.1	2.0
No. of probability words	M	9.2	5.8††	4.2††	4.4††	7.2
	S	6.1	5.1	5.0	3.9	5.0
	t		-2.9	-3.5	-4.0	-1.6
No. of different prob. words	M	6.8	2.1††	2.2††	2.5††	2.4††
	S	4.7	1.6	2.6	1.9	2.7
	t		-1.3	-5.1	-5.3	-5.1
H(CPT)	M	2.70	2.5	2.8	2.4	2.4†
	S	0.5	0.8	0.6	0.7	0.5
	t		-1.6	0.7	-1.7	-2.2
Percent correct for known questions		55%	53%	52%	59%	61%
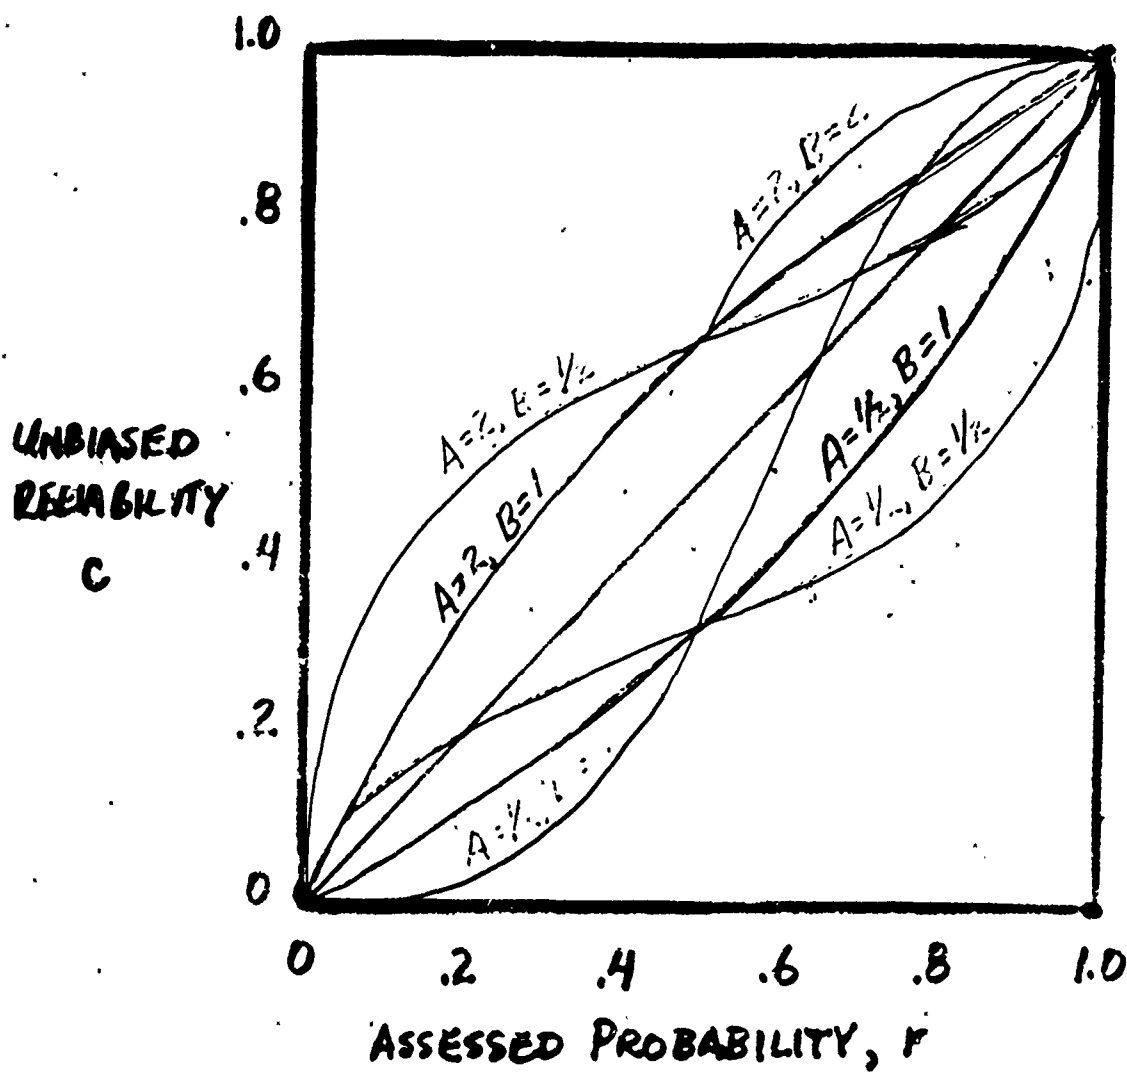
$*\ p(\mu_c > \mu_E) > .95$ $**\ p(\mu_c > \mu_E) > .99$

$\dagger\ p(\mu_c < \mu_E) > .95$ $\dagger\dagger\ p(\mu_c < \mu_E) > .99$

$$C = \frac{A\left(\frac{r}{1-r}\right)^B}{1 + A\left(\frac{r}{1-r}\right)^B}$$

Fig 2.3.

Fig 2.4

CALIBRATION CURVE
Hongkong students
N=53

Overall hit rate = 56.6%
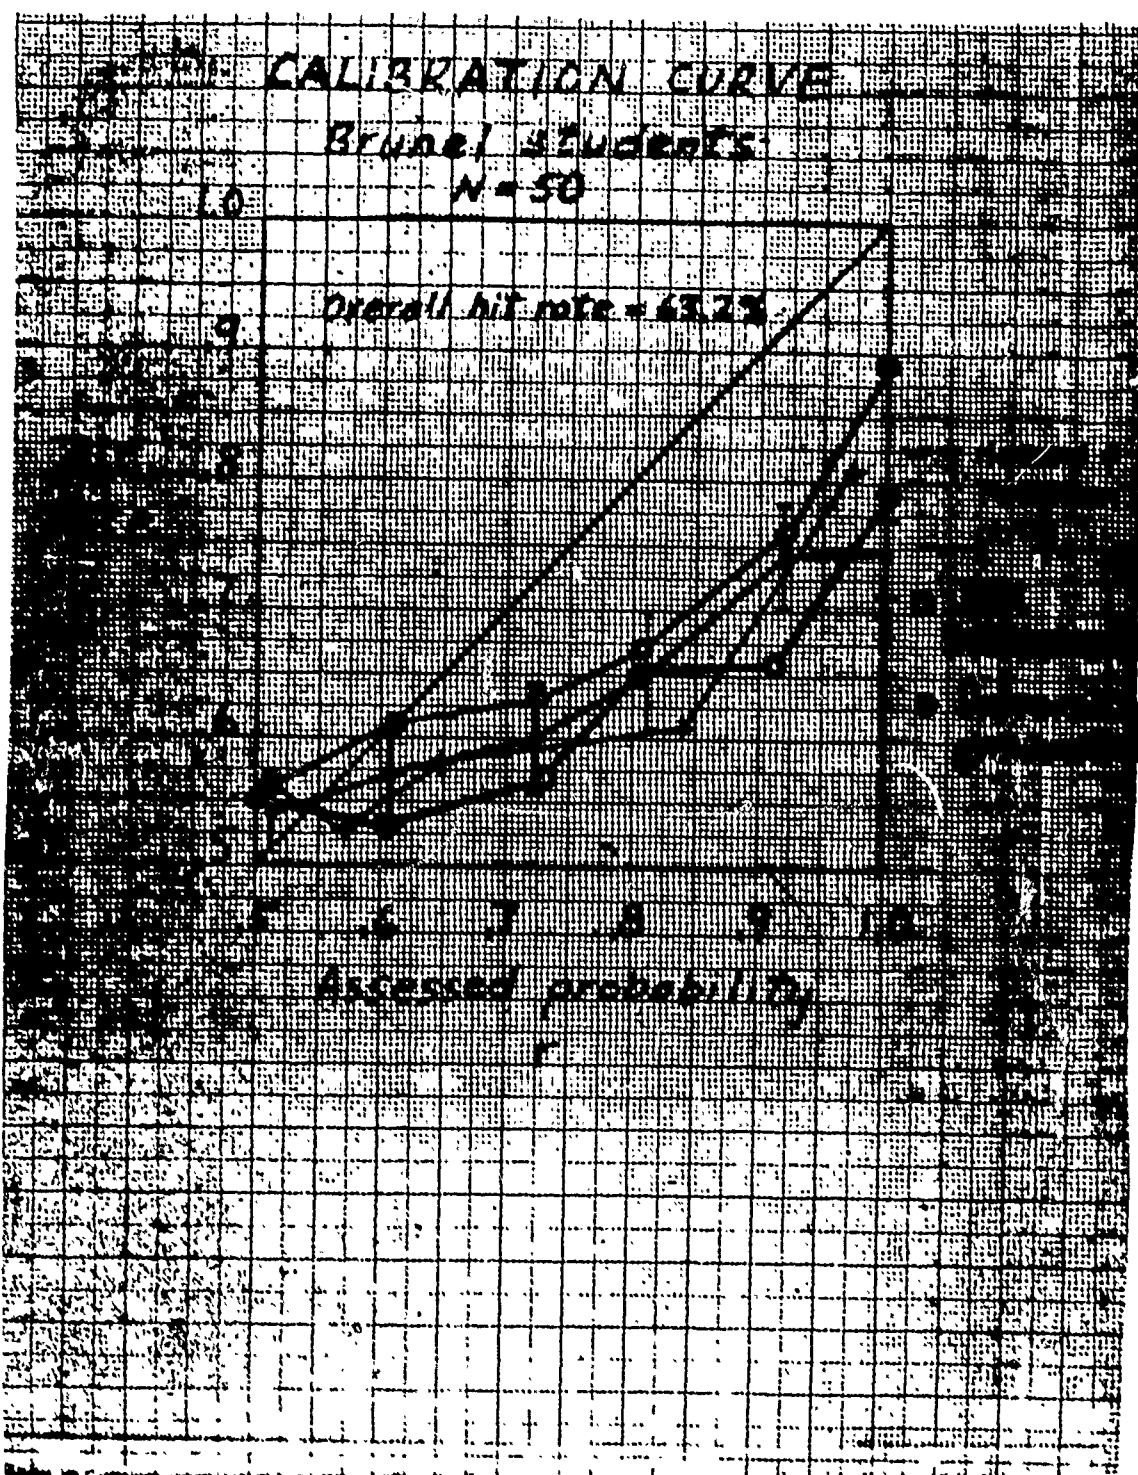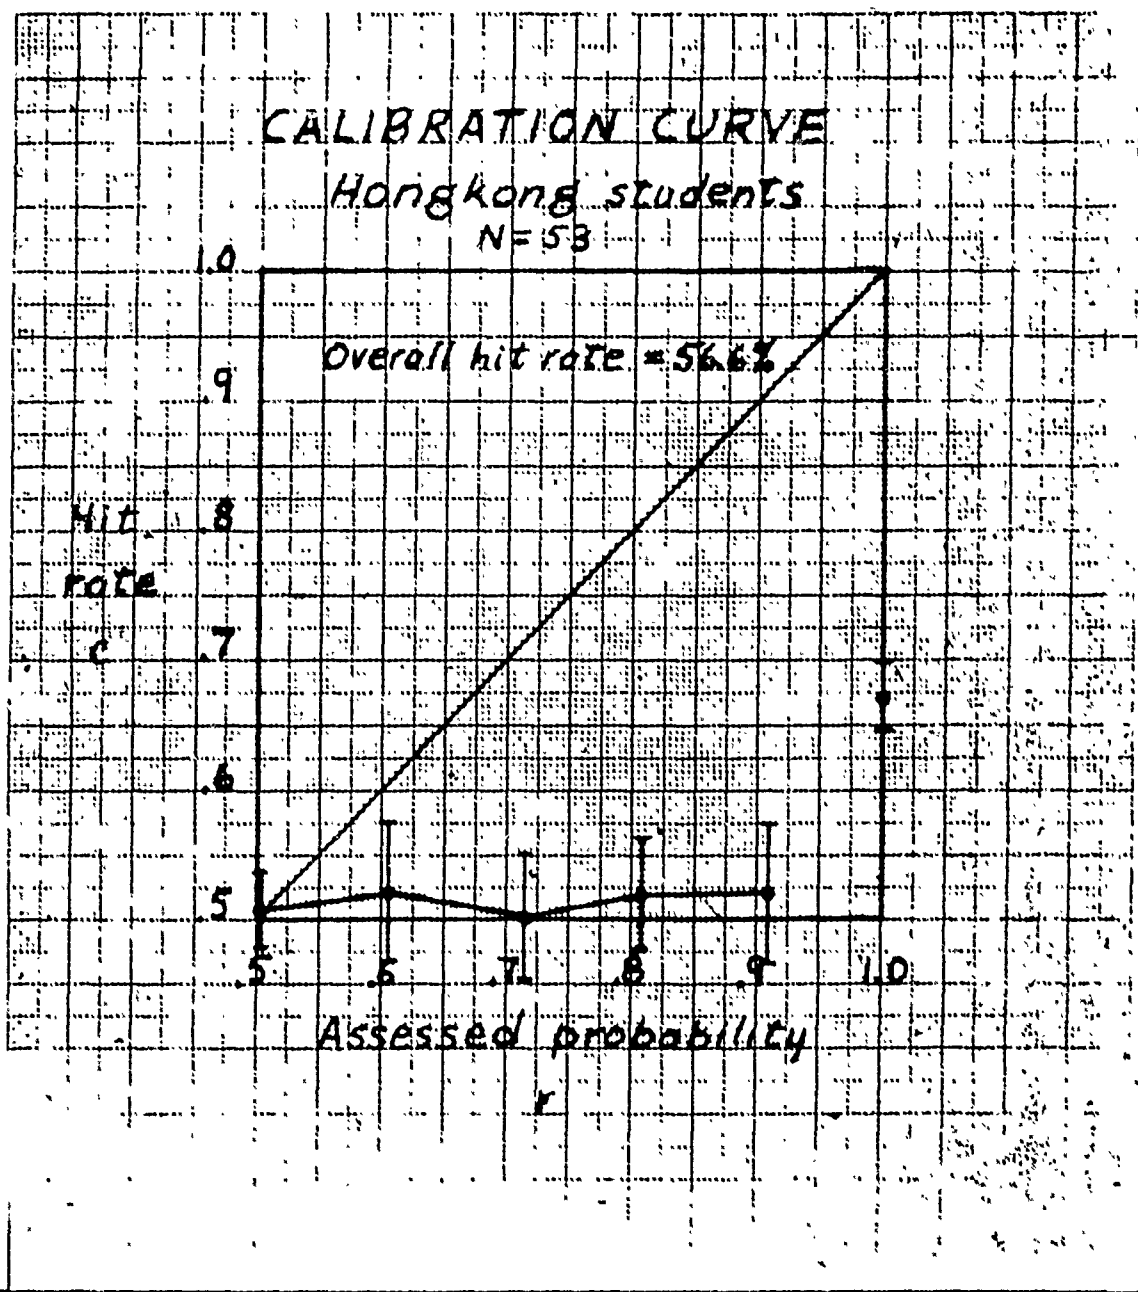
Fig 2.5

FACTORS

1	2	3	4	5	6
N.100 +	YN –	YN +	HR100s + [HR(NUB)]	A –
NC(ns) +	DK +	DK –		N50 +	B +
H[PAR] –		PN –		H (PAO) –	
		DPN –			

num.	verb.	verb.	num.	n/v	num.
numerically expressed disc. of uncert.	don't know	prob set. & verb exp of uncert.	knowledge ?	num exp of D/K	calib.

Fig 2.6

SESSION GROUP	No feedback 1	Feedback & training 2	No feedback 3	No feedback 4
A	PAQ 1	PAQ 2	PAQ 1	PAQ 3
B	PAQ 2	PAQ 1	PAQ 2	PAQ 3
C	PAQ 3	PAQ 3	PAQ 3	PAQ 1
D	PAQ 3	PAQ 3	PAQ 3	PAQ 2

PAQ = 75, Item Probability assessment Questionnaires
of which PAQ 1 and PAQ 2 are balanced verbal assessment
tasks.

PAQ 3 = Balanced 75 item geometrical task.

Fig 2.7. Design of Experiment

Session Group	1	2	3	4
A	5712	5405	5722	6368
B	5475	5586	5555	6234
C	6371	6706	6838	5242
D	6515	6748	6931	5439

Fig 2.8: Summary of Data - Quadratic Scores (Mean Scores)

"Cognitive Skills in Judgement and Decision Making"

Brehmer : I am going to talk about a specific topic: the role of
cognitive skills in judgement and decision making. But before I go
into this specific topic, I want to discuss the general framework in
which the topic fits.

The general framework is known as Social Judgement Theory (SJT).
SJT has grown out of Brunswik's probabilistic functionalism, and
it is a general framework for analysing judgement and decision
making.

SJT is committed to an analysis of the relations between two systems.
In our terminology, these systems are called "the task system" and
"the cognitive system". Their nature is illustrated in our general
model, the lens model (Fig. 2.9).

The lens model was developed by Brunswik, and it shows two systems,
extending in depth from a common boundary. This boundary consists
of the information available to a decision maker for his decisions
and is referred to as the "cues". In a medical decision problem,
the cues would be the symptoms exhibited by the patient.

To the left is the task system, defined in terms of the relations
between the cues and the outcomes, ie. the variable about which
decisions are made. This variable is called the "criterion". In a
medical problem, the criterion would be the disease of the patient.

To the right, we have the cognitive system, and this system is
defined in terms of the relations between the cues and the decisions
made by the decision maker. We are interested in analysing both of
these systems, but above all, we are interested in analysing the
relations between the systems. For that, we use a simple equation,
called the lens model equation. This is the mathematical representa-
tion of the lens model. The lens model equation describes the
relations between the two systems in terms of four parameters.

$$r_a = G \; R_e \; R_s$$

OUTCOME "CRITERION"   CUES   JUDGMENTS DECISIONS

TASK SYSTEM   COGNITIVE SYSTEM
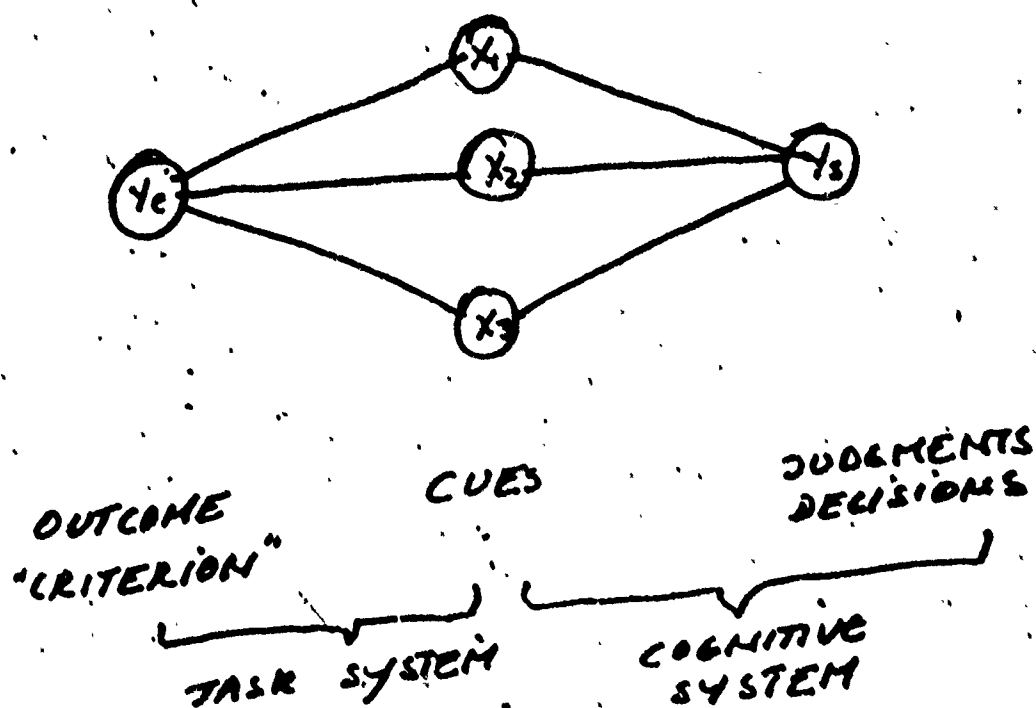
LENS MODEL EQUATION

$$r_a = G R_e R_s$$

$r_a = r_{Y_e Y_s}$ = CORRELATION BETWEEN CRITERION AND DECISIONS

$G = r_{\hat{Y}_e \hat{Y}_s}$ = CORRELATION BETWEEN PREDICTED SCORES IN THE TWO SYSTEMS

$R_e = r_{Y_e \cdot X_1 X_2 X_3}$ = MULTIPLE CORRELATION BETWEEN CUES AND CRITERION

$R_s = r_{Y_s \cdot X_1 X_2 X_3}$ = MULTIPLE CORRELATION BETWEEN CUES AND DECISIONS

Fig 2.9 The Lens Model   cont/..

$$\text{ACHIEVEMENT} = (\text{TASK PREDICTABILITY})$$
$$\times (\text{MATCHING OF SYSTEMATIC}$$
$$\text{FEATURES}) \times (\text{COGNITIVE}$$
$$\text{CONSISTENCY})$$

Fig 2.9 (cont/..)

SUBSTANTIVE ASPECTS

FORMAL ASPECTS

SURFACE CHARACTERISTICS

(1) NUMBER OF CUES

(2) METRIC CHARACTERISTICS

(3) INTERCORRELATIONS

SYSTEM CHARACTERISTICS

(4) THE RELATIVE IMPORTANCE OF CUES

(5) FUNCTIONAL RELATIONS

(6) ORGANIZING PRINCIPLE

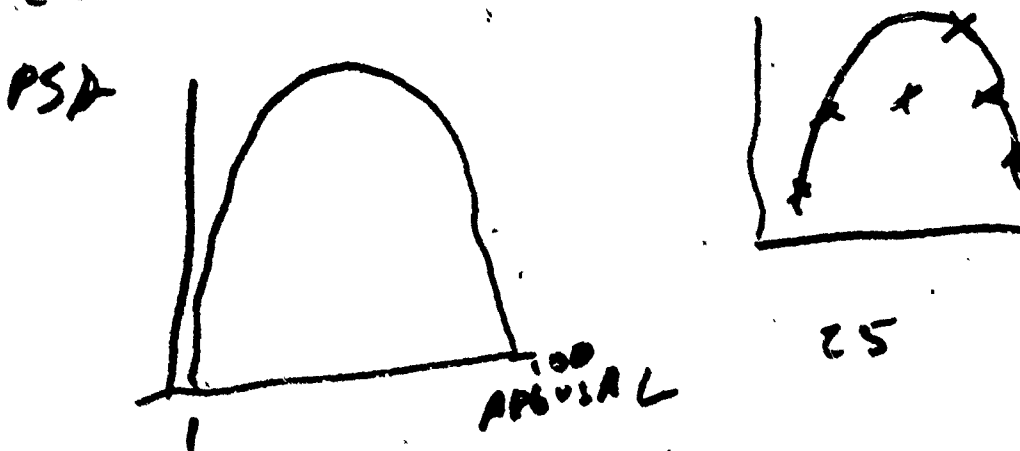(7) TOTAL PREDICTABILITY (R_e)

PSP

AROUSAL

25

Fig 2.9 (cont/..)

The first parameter, r_a, is simply the correlation between decisions and outcomes, and is a measure of the quality of the decisions. The second parameter, G, is the correlation between the linearly predictable variance in the task system and that in the cognitive system, and shows the extent to which the systematic features of the two systems match each other. The two multiple correlations, R_e, which is the multiple correlation between cues and criterion, and R_s, which is the multiple correlation between cues and decisions, show the predictability of the two systems.

The important thing about this equation for the present purposes is that it specifies the conditions under which a person will reach optimal performance. If the task system contains irreducible error, ie. if $R_e < 1.00$, decisions cannot be correct all of the time. The predictability of the task systems thus sets an upper limit for the correctness of the decisions. This limit will be reached if both G and R_s are one, ie. if the person both uses the information correctly, ie. gives the correct weights to the cues, and uses the information in a consistent way. G and R_s are statistically independent, except in the case when R_s is zero. In this case, G is automatically zero too. But for all positive values of R_s, G and R_s are independent.

To put the lens model equation into words, it means that the quality of decisions is dependent on the predictability of the task, the extent to which the information is used properly, and the extent to which it is used consistently. You might say that G expresses the decision makers knowledge about the task system, and R_s expresses his ability to use or apply this knowledge.

The fact that both of these aspects, knowledge and the use of knowledge, enter into the index of decision quality means that a cognitive system is always going to be specific to a given task. That is, some aspects of decision making performance just have to be acquired for each and every task, and this is knowledge about the structure of that particular task. If the decision maker does

not have this knowledge, no skills, abilities, heuristics or what
have you, will help. The G parameter thus shows that there are
limits to what you can achieve by general training of decision
skills, and the like.

We may now ask whether people are able to be optimal in the sense
defined by the lens model equation. The available results indicate
that people are able to be quite optimal with respect to G, ie. they
are able to learn the systematic aspects of most kinds of tasks, at
least within the area of clinical inference, which is where most of
the results come from. With respect to R_s, on the other hand, people
are not optimal, and we have yet to find a judge or decision maker
who is perfectly consistent in his utilisation of what he knows.
It is about this aspect of decision making performance that I am
going to talk, ie. I will talk about the use of knowledge, rather
than acquisition of knowledge.

The problem is whether you are able to apply what you know about a
decision task, for example, whether you are able to combine
information in your head according to whatever rule you want to use.

Before we go further into this, we will need to know what kinds of
demands are made by the task. Within SJT, we distinguish between
two aspects of a decision task: the substantive aspects and the
formal aspects. The substantive aspects refer to the content of
the task, ie. what specific things the task is about. The formal
aspects are a set of characteristics which may be used to classify
decision tasks without regard for the content. The characteristics
may be called the statistical aspects of the task, if you wish.

We distinguish between two kinds of formal aspects. The first kind
is called the surface characteristics. They are called the
surface characteristics because they pertain to the cues, ie. the
interface between the systems which is the only characteristics of
the task that can be directly observed by the decision maker when
he is making his decisions.

There are three kinds of surface characteristics.

The first is the <u>number of cues</u> that you have to take into account. The second is the <u>metric characteristics</u> of the cues; (the cues may have quantitative characteristics, or they may only have nominal characteristics). The third dimension is the <u>intercorrelations among the cues</u>, or the redundancy of the cues.

The next set of characteristics is called the system characteristics because they pertain to the relations between the cues and the criterion. There are four different system characteristics (namely, fourth to seventh).

The fourth is the <u>relative weights</u>, or relative importance of the cues; all cues may be equally important, or some cues may be more important than others.

The fifth characteristic has to do with the <u>forms of the functions</u> relating cues to criterion. These may be linear, so that the more you have of the cue, the more you will have of the criterion, or they may be nonlinear. For example, there may be an optimum value, so that middle cue values lead to high criterion values, and high and low cue values lead to low criterion values.

The sixth characteristic is called the <u>organising principle</u>, and refers to the rule to be used for combining information from different cues. This rule may be a simple additive rule, or it may be a configural rule, where the weight given to a cue depends on the value of another one.

The last (seventh) characteristic is the <u>total predictability</u> of the task, given the optimal organising principle. R_e is a measure of total predictability, but only when the task is linear and when the organising principle is an additive one. Otherwise, you will need different measures, or to transform your variables before R_e is computed. Total predictability is a property of the task system, or the cognitive system, given the available information. If new information is introduced, R_e may change.

The system characteristics have important effects on the consistency

of the cognitive system.

The first important parameter is the total predictability of the task, which has a large effect on R_s, and as the task system becomes unpredictable, so does the cognitive system. This is very easily demonstrated in the laboratory, but it seems to hold also outside the laboratory. For example, if you study clinicians making decisions from test data, you will find that the less predictable the task, ie. the lower the correlation between test and criterion, the lower their consistency, ie. the lower the correlation between the clinicians' decisions and the test. We do not know why this is the case, but the results seem stable enough.

Vlek : Is that a general conclusion? My impression is that clinical psychologists and psychologists with highly unpredictable judgement are sure of what they are doing.

Brehmer : Yes, they _feel_ certain about what they are doing, but if you actually calculate the predictability of their cognitive systems, eg. the extent to which they give the same diagnosis for the same case at different occasions, you will find that they are unreliable, and that they become less reliable when the task is highly uncertain. But subjective certainty, or confidence, is not the same as predictability.

The dimensions of particular interest in connection with cognitive skills are the fourth, fifth, and sixth, ie. the relative importance of the cues, the function forms relating cues and criterion, and the organising principle. These are of interest because they specify things that the decision makers have to do. Thus, the fact that some cues may be more important than others means that a decision maker will have to give them different weights when making his judgements, and to make his global judgements in such a way that the different weights are reflected in these judgements. The same is true about function forms: the decision maker must take the forms of the functions relating cues

and criterion into account and make his global decisions in such a way that they reflect these function forms, eg. so that he does not miss the optimal value, if the cue is nonlinear. The sixth characteristic also specifies a skill, in that it tells us that the decision maker has to organise the information in a certain way, and to put it together according to a given rule.

The question, then, is : do decision makers have these skills? Are they able to assign weights to cues, to use functional rules, and organising principles? Given that they know what to do, are they also able to do it? Or, to use a more extravagant formulation, are decision makers able to think the thoughts they want to think? With that kind of extravagant formulation you will, of course, suspect that the answer to this question is going to be negative. Otherwise there would be no point in this formulation.

At our laboratory we have done a series of experiments which show something about the nature of the skills involved. Though these experiments are but a beginning towards a complete inventory of cognitive skills, the results are nevertheless interesting and important in that they give information about what we may reasonably expect decision makers to do. Some results concern the level of skill, and other results pertain to the problem of whether or not these skills can be improved by training.

We will start by examining people's ability to use functional rules, because this area is where most of our experiments are. It is quite easy to do experiments in this area. All you have to do is to tell the subjects what rules to use, and then have them try to use the rules, and assess their success in using them. This may be done simply by fitting some kind of model to the subjects' responses and then study the differences between this model and the correct model implied by the rule you asked the subjects to use.

The general results of these studies are not surprising. You find that subjects are better at using linear rules than nonlinear rules,

and better at using positive linear rules than negative linear rules. The differences among rules are not small. If you ask a subject to use a second degree function, eg. a U-shaped function, the correct rule may account only for some 50% of the variance in his responses. If, on the other hand, you ask the person to use a linear rule, the correct rule will account for between 90% and 95% of the variance in his responses.

The proficiency with which people use functional rules is dependent on many aspects of the task. For example, if you give the cue information in numerical form, rather than graphical form, people are somewhat better at using functional rules, but the relations among rules do not change: nonlinear rules are still used less well than linear rules.

The number of cue values is another important aspect. If the cue has few different values, say five, subjects perform better than if the cue has many different values, say ten or twenty. This is not surprising if you think of man as a system with a limited capacity for processing information.

We now come to the problem of relative weights. To assess a person's ability to use information about relative weights, you first have to know what he understands by the concept of weight. If you tell a person that a piece of information is twice as important as another piece of information, what does this mean to him, if it means anything at all?

Within the regression framework, there are three possible interpretations of the concept of weight. First, it may be interpreted in terms of variance accounted for, or relative weights as defined by Hoffman. This seems to the interpretation of the weight concept used in decision theoretic analysis, for example, in multi-attribute utility work. A second possibility is to interpret the concept of weight in terms of beta weights or slopes, ie. the concept of weight would indicate the rate of change in the dependent variable, given a change in the independent variable, ie. the cue. A third possibility, which may be distinguished from the second only under certain conditions, is to interpret weights in terms of correlations.

If you want to investigate this problem, you may do so in two different ways. First, you may give a subject instructions about a task. For example, you may tell the subject to give certain weights to the cues, eg. that the weight for cue A should be twice that for cue B. Then you give him a series of cases, record his judgements, and analyse the results by fitting some kind of model which allows you to assess what interpretation of the weight concept corresponds most closely to the results.

The second possibility is to first obtain the judgements, then ask the subject about what weights he gave to the cues in the task. Then these subjectively reported weights can be compared to the weights obtained by analysing the actual judgements.

We have done both of these kinds of experiments, and the results indicate that at least one of the alternatives can be rejected: subjects clearly do not understand weights in terms of variance accounted for. In these experiments, we cannot, however, distinguish between an interpretation in terms of slopes and an interpretation in terms of correlations because the cues in the decision tasks were orthogonal, and in that case, slopes and correlations are the same thing.

It would, of course, be most reasonable if the subjects understood the concept of weight in terms of slopes, rather than in terms of correlations. However, we do have some preliminary evidence from a study using the second method, ie. reports about weights after the judgements, which used a task with intercorrelated cues. The results of this study suggests that an interpretation in terms of correlations actually accounts much better for the subjects' behaviour than an interpretation in terms of slopes. However, I have grave doubts about these results, because they do not make intuitive sense.

Kaplan : When you say "subjects" you mean subjects with a particular background?

Brehmer : In the work we do in Umea, we mainly use high school students who are the same age as American first year university students. This gives a sample of much more varied background than you get if you use only college students. So by people, I mean healthy, nineteen year old high school students with different backgrounds.

Vlek and Kidd : You fix the correlation between cues?

Brehmer : If you do not make the cues intercorrelated, you just cannot distinguish between slopes and correlations; they amount to the same thing.

Pask : You already cited the correlations and the redundancy of the cues?

Brehmer : Yes, we do have evidence from one study with inter-correlated cues, and this study indicates that the subjects understand weights in terms of correlations rather than slopes.

Johnson : Doesn't that differ from Sheperd's finding in the middle sixties?

Brehmer : Yes, but it would also be true that the cues that explain the greatest proportion of variance would be the cues with the greatest slopes in most cases. I do not think that Sheperd's study is critical to this question.

However, even though the results indicate that subjects understand the concept of weight in terms of slopes or correlations rather than variance accounted for, there is more to the problem than this. As you increase the ratio of the weights from a ratio of (say) 2 to 1, into 3 to 1, no definition of the weight concept seems to account very well for what the subjects do. Thus, it seems that we have the same problem here as with function forms: subjects

may not be able to actually apply the weights, even though they know what weights they want to use, and the more extreme the weight ratio, the greater the problems. These results may, of course, mean that the subjects really have no understanding of the concept of weight at all, but this would be contrary to the results of the first experiments, and therefore, I am more willing to believe that the latter results show the problems subjects have in using their knowledge.

Phillips : You mean that subjects have difficulty in understanding your instructions to use weights? This does not imply that they do not (intuitively) use them. It might mean they do not respond to the experimenter's request.

Brehmer : That is true, and that is why I want to talk about these results in terms of skills. To have a skill means that you can do things when you are asked to. If you ask a subject to give certain weight to the cues, and he then claims to have done exactly what you told him, and you find no interpretation of the concept of weight which reproduces what the subject has done, it seems reasonable to assume that we have a problem of lack of skill. Thus, I do not believe that subjects have a good internal understanding of weights, but refuse to show this understanding in our experiments; I think that they understand weights, but that they cannot show this understanding because of lack of skill.

Johnson : Do you give instructions or just give the training in the use of weights?

Brehmer : Well, in the studies I am talking about now, we only instruct the subjects about what we want them to do, we do not train them.

Phillips : Have you asked for the kinds of judgement that is typically Bayesian (for example, likelihood ratios)?

Brehmer : No, not yet, but it seems like an obvious thing to do, and we may try that.

I should say something about the organising principle too. So far, we have not done very much work in that.

There are at least two questions to be asked.

The first question is : do people use some organising principles better than others? For example, are they able to use a multiplicative principle better than an additive principle?

The second question is : what happens when people integrate information from different sources or cues?

With respect of the first question, we have not done very much. We have found, however, that people have a general tendency to integrate information additively, rather than in any other way, so, therefore, we use this rule in our experiments.

We have found, however, that it is harder for subjects to integrate information from a cue which is linearly related to the criterion with one that is nonlinearly related to the criterion, than to integrate information from two cues that are linearly related to the criterion. Specifically, we find that R_s, ie. consistency, is lower in the former case than in the latter.

With respect to the question of what happens when people integrate information, there are two separate problems which deserve study. First, you may ask whether the subjects are able to reproduce the components of the task when they integrate information. For example, if they have to integrate information from a cue that is linearly related to the criterion and a cue that has a U-shaped relation to the criterion, will their global judgements from both cues then be linearly related to the first cue and U-shapedly related to the second? The second problem is whether information integration leads to decreased consistency.

With respect of the first subproblem, an interesting possibility is
that information integration, which means an increase in cognitive
load, may force the subjects to simplify the task to be able to cope
with it. This may be done in various ways, such as ignoring some
information, or simplifying the task by, for example, making all
relations linear. Paul Slovic at Oregon Research Institute and I
have been investigating this problem in a series of experiments, but
so far the results have been negative in that subjects have been
able to reproduce the functions we have given them. That is,
integrating information does not necessarily lead to a simplification
of the task. The problem now is to investigate the limits of this
result, for example, by investigating the effects of time limits,
and the like.

With respect to the second aspect, ie. whether information
integration affects consistency, we have information that it does,
and that having to integrate information leads to lower consistency.
This suggests that information integration does indeed increase
cognitive strain, but apparently not so much that subjects have to
simplify their tasks.

To summarise this part of our results, they suggest that subjects
may not always be able to follow the rules for a given decision
task, even though they know perfectly well what they should do,
and that some rules are harder to follow than other rules. Clearly,
lack of appropriate cognitive skills is one of the reasons why
people do not perform optimally in decision tasks.

These results immediately lead to a question of whether these
cognitive skills can be improved by training. For example, is
it possible to improve people's ability to use functional rules
by training them in the use of these kinds of rules?

So far, all our learning studies have been concerned with the
problem of whether the ability to use functional rules can be
improved. The results are quickly told. Training leads to some
improvement, but not much, and we have not been able to find any
method for making people perfect in this respect. Furthermore,

the transfer effects are small. If you train a person with one
kind of rule, his ability to use another rule improves, but not
very much. What we get in our experiments looks more like a
general warm-up effect, rather than like genuine learning.

Phillips : Does this learning transfer?

Brehmer : Yes. If you train a person in the use of, say, an
inversely U-shaped function, and then test him with a new cue
variable with a different variance, he is able to use the function
as well as he was with the old variance. However, if you ask the
subjects to change their rule, eg. to make the inverse U-function
into a U-function, performance deteriorates.

Kaplan : Could you give briefly a specific example of training
for a task?

Brehmer : The best procedure for training that we could think of
was the following. First, we instruct our subjects very carefully
about what rule they should use. These instructions explain the
rule both verbally with examples, and graphically. Then we give
the subjects a series of cue values, and for each of these, they
try to come up with the correct criterion value according to the
function they had been shown. Subjects are asked to predict, say,
problem solving ability as a function of arousal, or something like
that, something that is reasonable to a psychology student, if that
is what your subject is.

Vlek : Their belief in the "reasonable" principle is, perhaps,
debateable.

Brehmer : Then you show them what the rule looks like, pointing to
the graph which illustrates how problem solving ability changes with
arousal. The subjects have all the information they need, as well
as examples of relations in their everyday lives which follow the
rules in question, so that you are quite certain that they understand
what they are supposed to do.

Kaplan : A short lecture, in fact?

Brehmer : Yes, and in addition to telling them, they receive the
written instructions. You do not give up until they know what to do.
Then you start giving them cue values, eg. arousal levels. They have,
of course, already been informed about the range of cue and criterion
values. After each prediction, they receive outcome feedback informing
them of the correct answer. In some experiments, we have used more
sophisticated forms of feedback, showing them not only the correct
answer, but plots of their own functions comparing these functions to
the correct functions after each block of ten trials. These experiments
were conducted with the hypothesis that since the subjects are supposed
to learn functions rather than responses, the kind of feedback they got
should be adequate to what they had to learn. Consequently, they should
be informed about their functions, rather than their responses.
However, this kind of feedback did not lead to greater improvement than
ordinary outcome feedback. This is somewhat surprising, I think.

Pask : Do they receive outcome feedback throughout?

Brehmer : Yes, but we have also done it without outcome feedback.
Quite a number of experiments have been performed, but they are mainly
just small variations on the general theme, and the general conclusions
are the same: subjects do not improve very much. Whatever improvement
takes place occurs very rapidly, say, in the first 50 trials, and it
does not pay to give your subjects many more than 50 trials; they just
get bored.

Pask : The case you have cited is for one cue?

Brehmer : We have, of course, also done experiments with more than
one cue. Using more cues just makes the problem more difficult for
the subject, but does not change the main conclusions.

Kaplan   : mathematically sophisticated people at an advantage (for
example, h a school students trained in algebraic operations)?

Brehmer : We have never correlated our results with mathematics
grades. We have, however, used third semester students in statistics
in one of our experiments, but they did not behave any better than
our ordinary kinds of subjects.

Kidd : What is the subjects' motivation?

Brehmer : We pay them for their work. Generally we do not feel that
motivation is a problem, but the subjects seem quite willing to
cooperate. Using paid subjects is quite different from using
psychology students as subjects who have to participate, and we
certainly prefer to use paid subjects.

Kidd : What is the difference between subjects who are forced to
participate and subjects who are paid?

Brehmer : We have never done any strict experiments on this question,
because it is simply not important to us. The only thing we can do
is to compare the results of an experiment with subjects who were
forced to participate with those of experiments where we use paid
subjects. The differences are not all that great, and you can get
psychology students to work quite well too, if you take good care of
them. For example, if you run subjects individually, in face to
face contact with the experimenter, most subjects work quite well.
However, if you run 100 subjects at a time, they may not work very hard.

Kidd : Do you think it would change the performance level if the
subjects found their experiment relevant to a real vocation or
occupation?

Brehmer : Well, I don't think so, but can only speculate on this
point. I do not think that it will not make much of a difference
because I think that the low levels of performance are due to cognitive
problems – limited processing capacity – rather than to motivational
problems. However, we have not done that particular experiment.

Pask : Motivation apart, there is also the question of risk attached to making real life (say, professional) mistakes.

Brehmer : The subjects in our experiments are not systematically wrong. It is not so much a question of systematic errors as a question of nonsystematic, or random errors. So it is not the case that the subjects feel that their judgements are going to be too low, or too high, all of the time. In that case, the subjects would probably give up. However, when you are making random errors, you will not have very good insight into your errors. In real life, you often are not informed about your errors, so you may not know that you have made a faulty diagnosis, for example, because the patient recovers anyhow.

Vlek : There is work in the multiple cue field that suggests you do not have to worry much about different kinds of functional relations but you do have to worry about the inter-relation of cues, since if the subject can choose cue decisions regarded as orthogonal a linear combination with unit weights is as good, or almost as good, as a more sophisticated scheme.

Brehmer : If we go back to the lens model equation, the Dawes and Corrigan argument has to do with G, and it says that almost regardless of what the task system looks like, provided the relations between cues and criterion are conditionally monotone, G will be high if the subject uses a linear additive integration rule. It would usually be better for a decision maker to use a unit weighting method, rather than any other method, because this method is easy to apply, and therefore R_s will probably be highest with this method. Using a straight unit weighting method without a human decision maker would, of course, be the best, because then you would have no problem with R_s at all.

Vlek : What happens if you instruct subjects to use a unit weighting rule?

68

<u>Brehmer</u> : I am sure that subjects could use a unit weighting rule
very well, but I doubt if they really act in this analytic manner.
Commonly, the descriptions people give of their judgement systems
are more complex than the equations you get when you investigate their
actual judgements. I think this is related to the problem of cognitive
skills. I think that at least part of the reason for the lack of fit
between subjective descriptions and actual equations based on
responses is lack of skills. If a decision maker knows what he wants
to do, but cannot do it because he lacks the skills required, and if
this lack of skills leads to a higher level of random error in the
cognitive system, a linear model is going to fit his responses quite
well, and better and better the higher the degree of randomness in
the responses, as is shown by Dawes and Corrigan. Thus, the more
complex a person wants to be, the less likely he is to succeed, and
the better a linear model will fit, and thus people will look simpler
when you look at their judgements, than when you listen to their
descriptions. But this is not because they do not know what they
want to do, or because they have no knowledge about the task, but
because they cannot do what they want to do. Therefore, I tend to
believe that you will get more valid descriptions and more useful
decision rules from people's verbal descriptions than from the
analysis of their responses. This problem needs more study, however.
It may be that the common practice of validating decomposition methods
by correlating the results from these methods with those obtained from
global judgements may not be such a good idea, because the global
judgements may not reflect what the person really wants to do.

<u>Vlek</u> : /Not verbatim/. Raised a number of questions. Much of
Dr Brehmer's arguments seemed to assume self chosen tasks; what
of tasks that are determined independently. In this case the
subject must <u>model</u> the given task (which is, itself, a cognitive
activity) and is bound up with the organising principles that are
used; subjects may be tacitly encouraged to build task <u>models</u> (as
internal cognitive representations) that are as "simple" as
possible: for example, being answerable to linear combinations
of unit weighting skills.

<u>Brehmer</u> :   As you know, Brunswik thought that you could have what
he called a psychology without a subject. He thought, then, that
you could go out and study, say, all decision tasks. I do not think
that this is possible, because decisions tasks depend on people, and
are defined by what people want to do. You have to study decision
tasks by means of subjects. But even then, it is not possible to
get a representative sample of decision problems before you do any
studies of decision making. To obtain a representative sample,
you first have to define the population of decision tasks, then
you can sample from it, but to define this population, you don't
need to know anything about decision making. Therefore, I think
that the inventory of decision tasks will progress as we learn more
about decision making, but it will not progress without studies on
decision making. As we learn more about decision making, our
conception of decision tasks will change. This is an answer to
your first question, I think. Furthermore, it is quite clear that
some of the most important decision tasks now are in people's heads.
Ken Hammond has remarked that we now live in a policy jungle, rather
than in the physical jungle. That is, what is important now is not
the actual going and looking for food, but our policy for producing
food. Thus, we have to involve people in our study of decision tasks.

As for your second question, that about unit weighting, I do not
think that the problem is to convince a decision maker that he has
to be more complex than he is. On the contrary, subjects almost
always describe themselves as more complex than their responses show
them to be, and possibly also more complex than the task requires
them to be, given the power of linear models.

<u>Pask</u> :   Your results, though rich, do depend (in their theoretical
interpretation) upon a lens model conception of decision making and
this leads, almost automatically, to statements (of the kind you just
made) about "how complex the task <u>requires</u> the response to be", or,
at a slightly earlier point, to the emphasis upon "simplicity"
and the like. Yet most complex decisions that are of consequence,
for instance a politician's, do not lend themselves to this type of
analysis and you seem to accept this, at least to some extent, in
commenting that subjects' descriptions are often more complex than
what subjects do as decision makers). Maybe what they do (cognitively,

or if suitably exteriorised, in behaviour) what they are up to is
"decision" in some sense of this word, perhaps "judgement".

One point, do you believe that complex tasks can be synthesised
from or decomposed usefully into, these elegant lens models; masses
of them if you wish. The other point, can we really see "decision"
or even "judgement" as an activity distinct from other mental
processes and beliefs. Perhaps the whole notion of 'decision'
and of 'judgement' is an artifact of how we, not the decision makers,
choose to image the task environment. These enquiries are not
perverse. They bear, in particular, upon decision training.

Brehmer : It is, of course, extremely hard to know exactly what
politicians do, and how much they do of this and how much they do
of that. But once you have started to use a certain framework for
analysing things, you tend to see examples that fit the framework
everywhere. This is, of course, not evidence that the framework is
true, or that it leads to true conclusions. It does, however, indicate
that the framework is useful, and we feel that our framework, ie, SJT,
can handle most of the problems we see in political decision making.
The framework also makes it very easy to communicate about decision
problems with people from other fields, eg. engineers or physicians;
they understand what you are talking about, and they can think of ways
of using our findings. I think that is an important aspect of SJT too.

Brehmer : We do a fair amount of work on subjects' ability to learn
judgement tasks. Indeed, this is our main research interest.

I think that you will have to make a distinction between the learning
of a judgement task in the sense of acquiring knowledge about that
particular task, eg. about what cues are relevant, what the relative
weights are, and so forth, on the one hand, and the problem of applying
what is known about the task on the other hand. It is quite clear
that subjects can acquire knowledge about judgement tasks, and that
they improve their judgements as a result of practice. This does not
mean that you can expect the subjects to be able to apply what they
have learned perfectly, however, and even though you may train a

person to do a task quite well, there may not be very much transfer
to another task because the basic skills are lacking. Consequently, I
do not think we have any discrepancy of results. We just have results
pertaining to different aspects of performance.

Phillips : For how long do you train these subjects?

Brehmer : We have given our subjects up to 800 trials.

Phillips : The difficulty is that this will modify the subjects
behaviour over several hours but over days or weeks cognitive
reorganisations take place which have little to do with strict
reinforcement theory. The situation is far more complex (for
complex judgements and task) than the literature usually suggests.

Pask : You propose the transfer effect may be due, in fact, to
simplicity of the situation.

A "Normative Process Tracing" Approach in Studying and Training
Decision Makers.

Goldsmith : I feel a bit diffident about presenting a method, or
an approach, that is still under development, but it seems to fit in
very well with the theme of the conference. For that reason I shall
present the approach but in a necessarily very general way.

One characteristic of this approach is, that it deals with complex
decisions; secondly, it attempts to define normativeness for a
complex situation; third, it is a descriptive approach, attempting
to describe what the subject is actually doing. It is process
oriented, it involves some decision analysis techniques, and in the
evaluation of the results it makes use of some principles from problem
·solving and cognition. The method or approach is concerned with the
transferability of training from one situation to another.

To begin with (Step 1) the experimenter chooses a decision problem
of sufficient complexity to interest the subject, and yet the problem
cannot be too complex so it looks unmanageable. Next, (Step 2) the
subject solves this decision task thinking aloud; a technique of
a process tracing where the subject endeavours to describe his thoughts,
strategies and so forth, while making his decision. Either before or
after a decision is reached (there are advantages and disadvantages
in each) the subject (Step 3) makes some numerical or graphical
assessments, as "hard" data, in contrast to the "soft" data provided
by the think aloud technique.

The purpose of all of this is to let the subject describe what he
was doing, not what upon more careful consideration, he thinks he
should have done.

When you have that descriptive data you go on to Step 4, acting as a
participant, to then apply a kind of decision analysis to the problem
that the subject has just dealt with. What you aim at here is to
attain an approximately normative model of what he should have done
(Step 2 and Step 3), so as to produce a decision task structure in as

close agreement as possible with the structure he would give it in
the light of his own perception of the problem and upon careful
consideration and his view of the world.

Next, at Step 5, you assess the effectiveness of the subjects'
strategies in attaining a structure in agreement with his semi
normative model and here your choice of methods must to some extent
depend upon the type of data that you've gotten from the think aloud
technique, and your strategy must be somewhat flexible. You can apply
such principles as the logical consistency or the adequacy of search,
the cognitive strain that was involved and the memory load which was
entailed by, or which you judge to have been involved in, the strategies
he used. Here you could seek the agreement of a set of experts to
evaluate the strategies from the standpoint of how effective they
were in leading to the semi normative model. There are two other
aspects you can evaluate also. It is possible to estimate the
realism or adequacy of the model (ie. the realism of his description
of what he was doing) and also to enquire how realistic is his
assessment of the effectiveness of what he was doing. I refer to
your attempt to assess the effectiveness of his strategies as
functional analysis; similar to an approach used in the study of
problem solving. How effective were his strategies in reaching the
goal, where he himself is to specify the goal in terms of the decision
analysis results you obtained.

Stage 6 is to devise a tentative training programme aimed at producing
more satisfactory strategies when he's tested again.

Step 7 is to go through stages 1 to 5 again to choose a new problem,
present him with that, study his think aloud strategies, obtain
assessments of what he thinks he was doing, how adequate he thinks
it was and to carry out a decision analysis on this new problem and
to assess the effectiveness of these strategies.

Finally, Step 8 is to take account of individual differences, and here
I would myself feel that the individual differences of relevance would
probably not lie in the area of conventional personality test results,
but rather in the performance of a subject in the situation itself.

I shall do no more than suggest that various studies have described
regular individual differences in think aloud behaviour, which have
been found to correlate with other aspects of decision making
behaviour or problem solving behaviour. Studies of the effectiveness
of a training approach designed for this particular kind of subject
involve a two by two by two factorial design; two types of subjects,
for each a training procedure which you held to be more effective,
and one you held to be less effective; half of the subjects would
receive the type of training judged to be more effective for them,
and half the type judged to be more effective for the other type of
person."

NOTE
Dr Goldsmith emphasised the difficulties of his method relating to
problem complexity, realism and the objectivity achieved as a result
of functional analysis in Stage 5 and Stage 6.

Goldsmith : I shall say a few words about the directions I think
training techniques might take. It seems to me that there are
certain elements in a number of different deviations from normative-
ness, which have been reported in the literature, all of which could
be described as the occurrence of focussed rather than divided
attention. There are general findings of business men ignoring
uncertainty, of non obvious alternatives being ignored, non obvious
events being ignored, of "as iffing", of Shepherds results (subjects
actually attending to fewer aspects of a problem situation than they
believe themselves to be). The results of various researchers,
including some of my own, indicate over estimation of the probabilities
of compound events, compared with the estimates of the probabilities of
their components. This could be interpreted, using a "Lens type"
scheme, in terms of "as-if", or "best guessed" model.

Vlek : What kind of approach have you made?

Goldsmith : My subjects were faced with judicial type decisions,
and were given evidence and asked to estimate the probabilities of
various events subordinate to the question of guilt, some of these
events were simpler, others more complex and composed of these simpler
events.

<u>Vlek</u> : Have you thought about interaction between utility and
probability? Utility in the estimation of probabilities. This was
probability of guilt I presume?

<u>Goldsmith</u> : No, but this is related to events subordinate to the
question of guilt. If the accused committed a crime he had to have
been at the scene of the crime, but being at the scene of the crime
does not render the accused guilty.

There are results of my own in this legal decision making area which
suggested that subjects are more normative in their estimates of the
probabilities of compound events and used approaches which were closer
to those of scientists in the sense of being more analytical in the
sense of naming more sub-hypotheses, subordinate to the overall question
of guilt, also being more critical of doubting of the evidence, and
also of using more active search behaviour.

There's another aspect of the legal decision making area which I'll
take up briefly, since it indicates the pitfalls that you can get into
if you are too certain in your expectations of models that the subject
should be using. I've tested law students, having them make prior and
a posteriori probability estimates on guilt, and likelihood ratio
estimates of the evidence and so forth, and also having them make
estimates, using a signal detection type scheme. I've done the same
thing with public prosecutors and with a few judges and defence
attornies. Whereas the law students who I tested were quite ready and
willing to give you estimates, the experienced lawyers (these were
people with 8 to 30 years of experience in court), were extremely
reluctant to make estimates of this sort. Not only that, but they
claimed that those were not the questions they were asking themselves.
One subject talked of judicial truth, saying that this was what they
tended to think of, to what extent does the evidence <u>prove</u> guilt, not
what's the probability of the accused <u>being</u> guilty. There is a
Bayesian bias that's present in almost all the decision theoretic
writings on legal decision making, that the task of the judge is to
decide the probability of the guilt of the accused, given the evidence.
There is an alternative approach which stems from some logicians and
legal scholars which, in essence, states that what the judge is asking

himself here is the probability that the evidence proves the thesis,
proof being defined as the existence of a causal relationship. That
seems to me too narrow, my own formulation would be that its the
probability that a causal or implicative relationship exists between
the evidence and the hypothesis, such that the evidence is indicative
of the thesis.

There's a lot to be said about this. The immediate reaction of all
these Bayesians (I, myself, am one of them) is that there must be
something wrong. Bayes theorem gives a means for calculating changes
in the probability of a given hypothesis given new evidence. Here,
however, the hypothesis changes and if you take new evidence, it is
no longer the same hypothesis.

Pask : What is the difference between "judicial truth" as you use it
there and "coherence truth" as Rescher would use this term. For
example, given propositions representing data, all of which are truth
candidates, we wish to estimate the truth of any one. In order to do
so implications between the propositions and their coherence with a
set of accepted propositions is taken into account (are the propositions
a logically coherent system) as well as their plausibility (one canon
of which may be past experience, another face validity and the like).

Goldsmith : Is the type of implication "formal implication" or "material
implication".

Pask : It could be either type. Presumably, in some systems of justice
formal would suffice but material and even causal implications are
possible. Does your criterion differ from Rescher's criterion?

Goldsmith : I would say that mine contains the element of cause,
that the critical relation is causal.

Pask : Isn't the logic of some legal schemes based upon a conventional
or traditional reflection of reality. For example, if a norm or even
precedent is taken into account, the evidence for accusation may rest
upon these factors as much as guilt.

Goldsmith : Those are factors that contribute to the estimation of the probability of a hypothesis entailing guilt. The effect of the system that I just mentioned would be to concentrate your attention upon those, perhaps few, items of evidence that seemed to be causally, or directly, related to the theme of guilt, and to consider other evidence as simply providing background information, allowing you to better estimate the probabilities of these premises or implicative relationships.

Pask : Similar criteria apply to command and control systems.

Goldsmith : What I want to emphasise is that if the decision analytic approach is used, there is the danger that you will miss the boat here and not end up with what is a truly normative model. In the particular case of the judicial decision making, if you fail to comprehend that what the subject is probably doing is to ask for a probability of truth and not the probability of guilt, you might end up with quite misleading results; in general, you do need to bear in mind the importance of perceiving just what the real goals of the decision maker are.

NOTE
There was lengthy discussion involving Goldsmith, Kidd, Phillips, Pask and Vlek.

Zeidner : As I understand it, the procedure that you followed, you relied very heavily on the verbal reports of your subjects? I was wondering if you had any problem in quantifying this information, establishing a structure, and what reliability you yourself feel in dealing with these kinds of reports, since you place heavy emphasis on verbal interaction, rather than observable behaviour.

Goldsmith : Of course, you could easily call these reports an +ʀ "observable behaviour", but be that as it may, I also use objective choice data.

Zeidner : But what happens when you are penetrating some more complex aspects· asking subjects why they thought this way and that way when they were rationalising their responses. I'm interested in how you handled, or how you would handle, this as an experimenter.

Goldsmith : There is a personal component there which you can't get around, as this is an ill defined situation, and that the experimenter must be as effective as possible to try to prevent the subject from just playing games. But how to define that component, I do not know.

Johnson : Did you do decision analysis? I guess I don't quite understand the role that decision analysis plays.

Goldsmith : To give you a normative model for what the subject has already done.

Pask : A personalised normative model?

Goldsmith : Personalised normative model, on the assumption that the subjects differ in their goals.

Johnson : Yes, yes, this is based on his verbal material.

Goldsmith : But the verbal material was then compared with this normative model that you got afterwards and, of course, carrying on a decision analysis you have to be verbal too, only that's not

the descriptive material or the behaviour that you get out of the think aloud technique prior to that.

Kleiter : I doubt that it is possible to make such general comparisons between decision analysis and protocols. In some investigations I tried to make an estimate of the number of stages that a subject is planning ahead in a multi-stage decision task using a suboptimal mathematical model for cross validation. We let people think aloud and just score the protocols for one variable (stages ahead); a very restricted approach but this is a terribly difficult step. I imagine you get into difficulty if you try to compare a complete decision analysis to the verbal protocols you have.

Zeidner : Dr Goldsmith, what concerns me is this (though perhaps
the analogy to your situation is not very good). Many years ago,
when I was interested in the problem of aerial interpretation as
detecting military significant objects from aerial photographs.
At that time we didn't have good methods of measuring eye fixation
or eye movements; as a matter of fact Mackworth had just started to
develop one. Some of our subjects were very good in doing this task
and we asked them, "do you do it systematically, do you judge by the
cultural features or the road network", and they gave verbal reports;
again in the context of selection and training. Later, when we finally
developed the methodology for eye movement responses, we used the
same individuals, there was no correspondence at all between a subject's
verbal reports, the reason that he gave for his identification,
and what his eyes were doing, ie. the objective evidence.

Phillips : I've been recently training underwriters in the use of
decision analysis, and one underwriter attended an extensive one
week course I gave at Brunel on decision analysis and went back and
invented some case studies. I kept discovering that the figures
that appeared at the ends of his decision trees were curious figures,
in the sense that they were always higher than the figures that I would
get when I would further decompose that event. For example, if you
think about the San Fransisco earthquake loss, what would the fire and
quake claims be? If you ask an underwriter you'll get a quite high
figure. If you decompose the situation by asking for a probability
distribution, and take the mean of that distribution, you'll always
get a lower figure. I began to discover that underwriters are using
some kind of heuristic which employs a concept they call "expected
maximum loss", or "probable maximum loss" (if you can stand the
contradiction in terms), and as far as I can make out, comparing the
decision analysis results (which is what you're suggesting), with
the intuitive judgements normally used, they are roughly taking the
95th percentile. They are not taking the mean, they are consistently
going to the very high end of the probability distribution. That has
some implications for training, because I've been trying to get them
to think in terms of expected monetary value rather than the worst
possible thing that can happen.

Vlek : It seems to me that the advantage of having people think aloud
when carrying out explicit analysis doesn't generalise over tasks.
Apparently, in a task which a subject can do almost automatically,
like scanning the ground for relevant objects (aerial photographs and
their interpretations) you might expect a great discrepancy between
his reports and his actual behaviour. Whereas in the situations
Goldsmith and Phillips are reporting, you can hardly do your task
properly without making things explicit. Higher mental processes are
implicated. If you have to do the task too quickly, or without proper
aid, you tend to rely upon or you intend to use all sorts of strategies to
simplify your task, even, or to play it safe.

Zeidner : I disagree with what you say on only one point, the task
of making these detections in identification is an extremely complex
cognitive task or very skilled task, and its not like riding a
bicycle. Subjects put pretty nearly the totality of their being
into it.

On a different theme, consider Method A and Method B of doing a task.
Suppose it is shown experimentally that Method B is just as effective
as Method A but subjects have been schooled in Method A, the school
solution is in Method A. Though you can show that Method B is more
efficient and show that there is no statistical or practical
difference in performance; they will still cling to the school
solution, even though its costly of time and effort. There is
tremendous resistance to change even in very simple tasks. I'll be
specific for a moment, again in the area of aerial surveillance.
People use stereoscopes generally to do this task because they feel
that the overlap, the 60^{o} overlap that provides the depth of dimension,
enables you to make better detections and identifications. Its a
very costly thing to do, and its very time consuming, and most of
the cues that people use do not depend upon the third dimension; all
this from results of simple, clear, experiments. In experimental
situations you get people with 19, 20 years of experience, and go through
the results with them at the end of the task, they say they have more
confidence in the response when they look at it stereo. But as
they're making each individual identification, as they're decomposing
the task, there is no difference in the degree of confidence.

I've grown increasingly sceptical, as you can see, about verbal reports. I'm also becoming increasingly sceptical about getting them to change their ways of doing things when they've been doing any task for a long time.

"Preliminaries For a Cognitive Decision Theory"

Vlek : " The history of this paper is that I became dissatisfied with
some restricted experimental work that I've been doing for some time.
I have exchanged the actual experimentation for some theoretical
thinking of which this is the preliminary result. My basic feeling
is that we can't go on doing experiments to see what variables
produce what kinds of behaviour without first trying to get an
overall view of the issues involved and should have at least a vague
notion of a cognitive decision theory. I have put the word
"preliminaries" in the title because I am certain that these are only
building stones for such a theory which will be very complex."

NOTE : Dr Vlek provided a revised version of his paper, taking account
of much of the discussion at the conference, and this version is
reproduced below. After it, there are segments of discussion, summarised
under specific topics so far as possible, which may add to the author's
point of view.

REVISED VERSION OF "PRELIMINARIES FOR A COGNITIVE DECISION THEORY" read
at the second conference on Decision Making in Compex Systems, Richmond
(London), 10-12 May 1976.

Abstract
Research on human decision behaviour should have two focal points :
(1) the development of a theoretical framework for the description of
cognitive functions involved, and stages and strategies occurring in
processing information for decisions; and
(2) the development of a functional classifications, or taxonomy, of
decision problems.
Ideally, one should like to be able to describe, or map, the latter in
terms of the former. A 'parent model' consisting of a sequence of
stages which might characterise purposeful behaviour is presented.
It is suggested that a full-fledged cognitive decision theory should
encompass such stages as, on one end, the recognition of a problem, and,
on the other end, the observation of, and learning from consequences

of a selected course of action. A taxonomy of decision problem
variables is proposed, and the consequences of connecting the
'parent model' to this taxonomy are discussed. Finally, a number
of hypotheses regarding Everyman's decision behaviour is listed,
and answers are provided to items of a questionnaire concerning
the training of decision makers.

Introduction

Pretheoretical considerations and empirically weak speculations
abound in this paper. The greater part of its message is rooted
in misgivings about the future results of decision making research,
given that investigators will continue to focus on narrow experimental
paradigms, fail to inventory the variety of Everyman's 'big decision'
problems, and refrain from attempting to make connections with other
areas of psychological research. Another motive underlying the paper
is the desire to develop a broad framework of thinking about adaptive,
purposeful behaviour based on judgemental information processing.
On the basis of such a framework the hypotheses and empirical results
of smaller-scale investigations might be more easily interpreted and
fitted together. Naturally, the risk of such a venture is that one
becomes the prime target of criticisms like Rabbitt's (1976) :
"This is the age of the easy psychologist. One by one the journals
capitulate to the easy writers of our time. Papers expand to
monograph length. Introductions and discussions sprawl and lay
claims to more than their authors know. Experiments are performed
as excuses for speculative harangues rather than to test points
of fact. Theories are held of more account than evidence, and
hypotheses are valued less for the precision with which they
are worked out than for their contradictions of uninformed
intuitions". It is hoped, however, that the tolerance
level for speculation at an informal conference like this, is
sufficiently high to make this contribution acceptable, and
Rabbitt's criticism less applicable.

Also,the present ideas should be held against the background of a wealth of theoretical and empirical studies of decision behaviour,in order to document their validity and fruitfulness for further research.This meticulous check against existing literature is the subject of later work.

In the remainder,five sections are distinguished.First,I should like to make a few rather primitive statements concerning psychological functions involved in human purposeful behaviour.Secondly,a scheme of hypothetical stages stressing the judgemental information processing aspects of such behaviour is presented.Third, 'it will be argued that such a scheme of information processing stages applies differently to different decision problems;hence a meaningful classification of such problems is needed,and will be proposed.Fourth,a number of hypotheses regarding overyday human decision behaviour will be listed.And fifth,answers will be given to a questionnaire put before the participants of this meeting.

1.Psychological functions involved in human decision behaviour.

The basic question of this conference is:can individuals be trained to be effective decision makers? Initially,that provoked an immediate 'no', because it seems that decision making pervades human behaviour to such an extent that almost every human function is involved.So in fact one could say that effective decision makers are effective survivors in this world,and that therefore training should cover the basic survival skills.Of course,there are simple survival problems for which only a few survival skills are needed,whereas complex problems call for many skills.This is illustrated in table 1,where along the rows you find some choice problems increasing in complexity from top to bottom.Along the columns you find some areas of human functioning.Of course these do not form a proper categorisation of human functions,but they have at least been the topic of psychological research (except perhaps 'imagination',which is understood to be a sort of 'cognitive simulation' of future actions,ovents,and consequences).

decision problem:	explicit and conscious use of function:						
	perception	learning	memory	imagination	evaluation	choice	action
skin irritation	-	-	-	-	⌐	-	x
coffee vs tea	-	-	x	-	-	x	x
2 known menus	-	-	x	-	x	x	x
Paris vs London	-	-	-	x	x	x	x
Jobs	x	-	x	x	x	x	x
Houses	x	x	x	x	x	x	x
Investments	x	x	x	x	x	x	x

Table 1: psychological functions involved in simple and complex decision making. x-marks indicate that a function is explicitly needed;dashes indicate functions playing a secondary role.

Consider the rows of table 1.When we become aware of skin irritation we usually don't do anything more than just (re)act:we scratch our skin.'Coffee versus tea'. requires something more:you have to (very quickly) remember what they are,choose, then drink it.When deciding between two known menus we usually try to remember what they were like recently,then evaluate,choose,and order one of them.Deciding between Paris and London for a holiday trip by a non-European requires imagination (if he has never been there),evaluation,and choice,before the trip can be started. When it comes to choosing between jobs one has to perceive carefully what is being offered. One may try to recall aspects of previous, or other people's, jobs, and the last four column functions apply as before. In the choice between houses or investments all listed functions come into play, since such kinds of decisions may improve by experience. What table 1 illustrates and suggests is that <u>a greater complex of human functions is involved when one is dealing with more complex decision problems.</u>

What is generally called 'decision theory' has in fact two branches.One is descriptive-psychological;it concentrates on charcteristics of human decision behaviour.The other is normative-technical;it provides rational procedures for the optimal solution of decision problems. I prefer to consider the normative-technical branch to be a smaller branch of the main,descriptive-psychological branch: any normative application of decision theory is necessarily and inevitably embedded in a larger (socio-psychological?) structure which cannot be organised in a normative way (except perhaps under dictatorship).This larger structure should be described in order to reveal the psychological context of the normative analysis.It is essential that great care is taken in establishing the limits of applicability of normative decision theory.Which actions are considered in a decision analysis and what their consequences are believed to be depends greatly on the decision maker's theory of the task environment.Such a subjective,individual theory summarises the decision maker's view of essential factors or mechanisms underlying a given problem situation.Only an effective knowledge of the basic structure of a task environment will enable the decision maker to establish with some certainty what is likely to happen as a result of his and his opponents' actions.This,I think,is the meaning of the term 'expert knowledge' about a situation,which - in a much more restricted fashion - returns as the concept of 'substantive goodness' in probability assessment.Such functions as listed in the first four columns of table 1 play an essential role in the formulation of subjective task environment theories.

2. A tentative scheme of stages in (complex) decision behaviour

From a decision-theoretic point of view purposeful behaviour is aimed at maintenance or increase of the   degree of satisfaction characterising the situation in which one has to live. We shall not here discuss the matter of short- versus long-term satisfaction, or present versus future satisfaction. These are aspects of certain kinds of decision problems (and decision makers) to which attention will be given in section 3.

For purposeful behaviour to occur, I suppose that people find themselves in some existing situation, which may be described in terms of utility dimensions, and of which they attempt to maintain or increase the aggregate utility by avoiding or correcting imminent deteriorations, or seeking improvements. In order to do so one often has to plan and undertake actions, and learn from experience.

This whole process can be divided up into a number of hypothetical stages with interconnecting loops through which one may go back as desired in order to 're-run' or revise parts of the process. Such a scheme may reveal psychological issues involved, and it may suggest points at which training could be most effective. Table 2 and its accompanying list of Clarifications largely speak for themselves. They concisely summarise what should have been the next several pages of discourse.

Table 2:a scheme of hypothetical stages characterising (complex) decision
behaviour (see Clarifications on following pages).

stage	recourse	stage description

A. problem recognition

A. 1) something desirable   }(certain
   2) a deteriorating status quo } or only
                                      probable.

1. "does this action solve my problem?"

B. planning of,& search for feasible actions

B. specification of goals ⟶ utility aspects ⟶ actions

2. "couldn't I find any more/better ones?"

C. listing of feasible actions

C. result of a preliminary evaluation: superficially acceptable actions

3. sensitivity analysis: "some actions have very similar consequences"

D. exploration of consequences & of possibilities for later action

D. — setting a time horizon (floating?)
— description of uncertain factors
— description of consequences
— description of contingent actions

4. sensitivity analysis: "some consequences are trivial"

E. assessment of utilities & quantification of evidence

E. — choosing a MAUT-model
— choosing an information integration rule

5. "intuitively, I don't agree with the computed rank order"

F. over-all evaluation of & choice among listed actions

F. — choice of a decision criterion

6. hesitation:"this seems best to do,but..."

G. initiation of chosen action

G. irreversible allocation of resources

7. cognitive dissonance: "I should (not) have ..."

H. subdecisions on adaptive/corrective actions

H. allocation of reserve resources:manipulating the task environment so as to maximize utility of consequences (they rarely occur exactly as anticipated)

8. "how can I get this back to normal?"

I. observation of,& learning from actual consequences of chosen action

I. — keeping statistics
— relating (remote) consequences to initial actions
— post-hoc evaluation
— emotional reactions

(next cycle)

CLARIFICATION OF STAGES AND LOOPS IN TABLE 2.

stage A : perception of status quo and of desired state;description of causal
relationships among elements of task environment;distinguishing super-
ficial symptoms from underlying factors;predicting future developments
of status quo;noticing relevant discrepancies between what exists and
what is desirable

stage B : specifying goals;formulating coherent goal-hierarchies;(well-specified
and detailed goals are suggestive of concrete actions);creatively
designing courses of action;recognising choice-points;(limitations of
attention;directedness of attention;traps of habitual thinking)

loop 1 : once the problem is defined,the relevance of proposed actions can be
judged; or:contrived actions may alter the formulation of the problem;
(possibility of skipping stage A altogether)

stage C : some actions are kept for further exploration;ability to judge proposed
action as (un)feasible,so that the remainder of the process deals with
actions "worth considering";ability to make quick over-all evaluations
of actions against standards of 'feasibility'

loop 2 : will be used as long as listed actions are too few in number

stage D : developing an (internal) decision tree through a process of cognitive
simulation (Toda 1976); ability to imagine remote consequences; ability to
predict changes in consequences as a result of future adaptive actions;
recognition of relevant contingencies

loop 3 : adaptation of list of feasible actions;deletion of dominated actions;
deletion of actions with consequences similar to those of other actions

stage E : listing of (independently contributing) utility attributes;setting
minimum goals in terms of utility attributes;weighing relative
importance of utility attributes;being aware of time preferences and
possible discounting of remote risks;overcoming judgmental biases in
scaling utility values and importance weights;identifying independent
pieces of evidence about uncertain events;contriving a scenario for
complex uncertain events;integrating evidence into unbiased probability
statements;willingness to make one's (true) beliefs and values explicit;
dangers of interacting probability and utility (optimism/pessimism)

loop 4 : may lead to pruning of the chance forks of the (internal) decision
tree;recognising irrelevant uncertain events as having similar conse-
quences as their outcomes;deleting branches leading to zero utility

stage F : integrating beliefs and values about consequences into an over-all
estimate of a course of action's attractiveness;testing estimated
utilities against set minimum levels;discriminating degrees of attract-
iveness;picking out the best possible action;willingness to choose
randomly among equally-valued actions

loop 5 : intuitive check on results of the procedure so far;accepting decompos-
ition-and-recomposition as a rational procedure;overcoming the resistance
of the (unaided) 'professional eye'

stage G : "getting the thing going";overcoming desires to extend the decision
process and delay action;accepting course of events and possible actions
contingent upon current choice;trusting one's potential to adapt to
undesirable consequences;willingness to make irriversible investments;
possibility of wishful thinking

loop 6 : self-evident (see stage G);one should try to keep out of it,though using
it may increase one's confidence in the chosen action

stage H : choices at points left open when the chosen action was designed;
reactions to un-anticipated side-problems

loop 7 : (one cannot reverse what happened at stage G;therefore loop 7 reaches
back to stage F) ; feelings of regret,or post-hoc rationalisation may
distort memory of the preceding decision process

stage I : (stages H and I are in fact alternating); possibilities for probability
and utility learning;most critical:recognising the immediate/remote
consequences of one's own actions;emotional reactions,e.g. the tendency
to undo decisions having resulted in very bad consequences (attempting
to simulate the more desirable which did not happen)

loop 8 : ongoing cognitive dissonance phenomena (reaching back to stage F);
planning of adaptive actions (back to stage H);possibility of panic-
or reckless reactions to extremely bad or pleasant consequences,resp.

- - - -- - - - - - - -

Of course,the stages in table 2 are global.Each one of them may be further
broken down into substages,whose nature may depend on particular details of the
decision problem at hand.Toda (1976) uses the metaphor of a neutral observer,
who may either stand at great distance to survey what is going on,or come very close
by with a magnifying glass to inspect the fine details of the process.Examples
of a fine-grained analysis of what should happen at stages A-D may be found in
the decision analysis literature (e.g. Howard et al,1976).Similar examples
pertaining to stages E and F occur in the psychological literature on decision
behaviour (e.g. Payne,1976;Montgomery & Svenson,1976).

It is important to notice that fine-grained analyses of what happens at other
stages or between most stages may be found in other areas of psychological
research;but their results so far have not been integrated into a common frame-
work.

3.Toward a classification of decision problems

In the foregoing,the terms 'purposeful behaviour' and '(complex) decision
behaviour' have been used more or less interchangeably.This apparent sloppiness
results from the possibility to consider the scheme of table 2 as a 'parent model'
of behaviour,if the latter is studied from a general decision-theoretic point
of view.Extremely 'big' decisions calling for explicit consideration of a multi-
tude of aspects of the situation would mobilize all stages and loops of the

parent model.Much simpler choice problems (cf.table 1) would,of course,
allow the decision maker to skip several stages because he may take their
result for granted.Further development of the scheme in table 2,and subsequent
adaptive use of it should create possibilities to account for behaviour varying
in the range 'between decision and habit'.

While using the parent model of table 2 in this way we shall soon notice
that the actual course of a particular decision process is heavily dependent
upon the nature of the decision problem.It may be either very clear or unclear
what the decision maker aspires to achieve.Often there are only two options
each of which has probable serious consequences in the long run,which are
difficult to foresee.Or,consequences may be short-term but confusing by their
multidimensionality.Sometimes the sheer number of alternative courses of
action prohibits a thorough exploration of the possible consequences of
each single one.Or,the decision maker does not even know for sure whether he
will be able to carry out the selected course of action.

.People attempting to develop a task taxonomy meet the difficulty of choosing
a conceptual foundation.What is a task? It may be useful at this point to
paraphrase Fleishman (1975) who subscribes to the statement:"A task is any
set of activities,occurring at about the same time,sharing some common
purpose that is recognised by the task performer".Fleishman does not pretend
to offer the definition of 'task',but discusses four conceptual bases for
classifying tasks:(1) behaviour actually occurring while the task is performed,
(2) behaviour requirements for proper task fulfilment,(3) required abilities
for proper task fulfilment (more fundamental than (2)),and (4) task charact-
eristics as intrinsic,objective properties of the situation (goal,material,
procedure,response-devices,etc.).

The scheme of table 2 would seem to offer possibilities for classifying
decision behaviour in terms of the conceptual bases 1,2,and/or 3 just mentioned.
The taxonomy of variables characterising decision problems,given below,on the
other hand,rests on intrinsic properties of the problem,its possible solution
alternatives and their consequences,etc.,as well as on observable
        properties of the decision maker.

Table 3 presents building elements for such a taxonomy of decision problems.
Again,the contents of table 3 should speak largely for themselves;discussing
and illustrating the several variables one by one would require considerably
more time and space.It is hoped that the reader has sufficient knowledge of
the literature on decision making to be able to grasp the meaning of the
various lines.

parent model.Much simpler choice problems (cf.table 1) would,of course,
allow the decision maker to skip several stages because he may take their
result for granted.Further development of the scheme in table 2,and subsequent
adaptive use of it should create possibilities to account for behaviour varying
in the range 'between decision and habit'.

While using the parent model of table 2 in this way we shall soon notice
that the actual course of a particular decision process is heavily dependent
upon the nature of the decision problem.It may be either very clear or unclear
what the decision maker aspires to achieve.Often there are only two options
each of which has probable serious consequences in the long run,which are
difficult to foresee.Or,consequences may be short-term but confusing by their
multidimensionality.Sometimes the sheer number of alternative courses of
action prohibits a thorough exploration of the possible consequences of
each single one.Or,the decision maker does not even know for sure whether he
will be able to carry out the selected course of action.

People attempting to develop a task taxonomy meet the difficulty of choosing
a conceptual foundation.What is a task? It may be useful at this point to
paraphrase Fleishman (1975) who subscribes to the statement:"A task is any
set of activities,occurring at about the same time,sharing some common
purpose that is recognised by the task performer".Fleishman does not pretend
to offer the definition of 'task',but discusses four conceptual bases for
classifying tasks:(1) behaviour actually occurring while the task is performed,
(2) behaviour requirements for proper task fulfilment,(3) required abilities
for proper task fulfilment (more fundamental than (2)),and (4) task charact-
eristics as intrinsic,objective properties of the situation (goal,material,
procedure,response-devices,etc.).

The scheme of table 2 would seem to offer possibilities for classifying
decision behaviour in terms of the conceptual bases 1,2,and/or 3 just mentioned.
The taxonomy of variables characterising decision problems,given below,on the
other hand,rests on intrinsic properties of the problem,its possible solution
alternatives and their consequences,etc.,as well as on observable
properties of the decision maker.

Table 3 presents building elements for such a taxonomy of decision problems.
Again,the contents of table 3 should speak largely for themselves;discussing
and illustrating the several variables one by one would require considerably
more time and space.It is hoped that the reader has sufficient knowledge of
the literature on decision making to be able to grasp the meaning of the
various lines.

Table 3:Taxonomy of variables characterising decision problems

1.Basic properties of the task

1.1.Nature and size of discrepancy between existing and desired states;
1.2. Extent of time during which chosen action (should) lead(s) to
     inevitable consequences;
1.3. (expected) Total number of possible choice alternatives;
1.4. Simultaneous and full availability of choice alternatives versus
     nonsimultaneous and/or partial availability;

2. Properties of choice alternatives

2.1. Certain vs uncertain consequences (risky vs riskless);
2.2. Dimensionality of consequences (single- vs multi-dimensional);
   2.2.1.)
   2.2.2.)   Varieties of multi-attribute utility:additive,lexicographic,etc.;
2.3. Reversibility vs irreversibility of consequences;
2.4. Time variability of consequences;
2.5. Single- vs multi-stage actions (multi-:later,adaptive actions possible);
   2.5.1. Sequential choices with changing beliefs and/or values;
   2.5.2. Dynamic choices among contingent sets of options;

3. Availability of choice alternatives (if nonsimultaneous and/or partial).

3.1. The n most recently considered alternatives are available for choice.
     With n=1,only the last-considered alternative may be chosen or rejected.
     With n=all,the DM may choose from an expanding choice set.
3.2. Partial availability:probability that the chosen alternative can be had
     (alternatives may 'refuse' to be chosen;'memory-stored' alternatives may
     disappear from the set;a chosen action may appear to be unrealistic);
3.3. Values of alternatives may result from random sampling from a generating
     distribution;
   3.3.1. Variations in shape and parameter values;
   3.3.2. Distribution characteristics known vs unknown to DM;
   3.3.3. Distribution characteristics stationary vs nonstationary in time;
3.4. Cost of acquiring the availability of a new alternative;
   3.4.1. Material vs 'psychic' costs;
   3.4.2. Uncertainty about costs;
   3.4.3.(Non-)stationarity of costs across sequence of alternatives;

4. Stress and other situational factors

4.1. Available decision time;
4.2. Multiple vs single task performance;
4.3. Noise,heat,fatigue,drugs,etc.;
4.4. Dominant presence of others;
4.5. Imposed memory conditions (especially:availability of an external memory);
4.6. Availability of information sources;
4.7. Availability of external resources to implement a chosen action;

5. Characteristics of the decision-maker.

5.1. Cognitive abilities;
   5.1.1. Memory,attention,abstraction,judgemental efficiency,etc.;
   5.1.2. Stage of psychological development;
   5.1.3. Amount of training and experience (decision procedures;task environment)
   5.1.4. Imagination:capacity to simulate contingent futures;
5.2. Motivational aspects;
   5.2.1. Extent of maximising tendency (aspirations);
   5.2.2. Patience vs impatience (time preferences);
   5.2.3. Tolerance for uncertainty/ambiguity;
5.3. Personal mental and physical resources for implementation of chosen actions.
- - - - - - - - - - - -

Connecting the taxonomy of table 3 with the parent model of decision
behaviour given in table 2,creates the interesting job of adapting the parent
model to the several kinds of decision problems following from table 3.
Ideally one should like to be able to describe each category of decision
problems in terms of the nature and ordering of psychological functions,stages
and strategies involved in solving problems of that category.Each attempt to
map a decision problem category into the scheme of table 2 provides a test
for the validity of the latter.Of course,the whole enterprise of connecting
a taxonomy of decision problems to a parent model of decision behaviour will
neither be easy nor quickly be finished,and it is possible that the parent model
may have to undergo drastic revisions.A coupling of the two,however,seems very
necessary for any parent model of decision behaviour and any taxonomy of
decision problems,given the likelihood of a strong task dependence of human
cognitive processes.

4.Some hypotheses about human decision behaviour

Let us,to further stimulate the discussion,consider a number of hypotheses
concerning decision behaviour.Some of these are old,some new,but most of them
pertain to everyday decision making,and they have (therefore?) been relatively
neglected by students of decision making.It is likely that empirical evidence
concerning these hypotheses is buried in diverse areas of psychological literature.
This information should be uncovered and brought into contact with the
decision-theoretic point of view about behaviour.Undoubtedly,interesting
re-interpretations will result if such is done.

Thirteen hypotheses follow here without further comment.

a) Decision makers like the status quo,because of all possible courses of action
it permits the most detailed cognitive simulation of the future;

b) Emotional reactions to discrepancies between existing and desired states
blur one's capacity to distinguish action-relevant from action-irrelevant
discrepancies (about which one cannot do anything);

c) Time stress on search for feasible actions results in a listing of short-term
problem solutions,since cognitive simulation is easier on the short than on
the long term;

d) Time stress on evaluation of possible actions leads to concentration on key
utility attributes;

e) Subjective time horizons,both forward and backward,depend upon the subjective
value of (experienced or anticipated) consequences;

f) In judging values and probabilities uncertainty about the proper estimate
will produce 'regression towards the middle of the scale';

g) Judgements of values and probabilities are subject to biases
elicited by the available response-device;

h) If available resources are scarce, exploration of alternative
actions may be restricted to one or two which meet minimum require-
ments on key utility attributes;

i) High probability/low damage risks are considered (and avoided)
more carefully than low probability/high damage risks;

j) Damaging consequences provoke, when they occur, a tendency to deny
them and to simulate their positive counterparts instead;

k) Illusions of personal responsibility lead to refusal to flip a coin
in case two actions are judged equally attractive;

l) Information about choice alternatives is processed sequentially
rather than simultaneously; in combination with selectivity of attention
and limitations of memory this is an important cause of irrational
behaviour; it also makes 'satisficing' a more appropriate decision
strategy than 'maximising';

m) Exaggeration in the daily use of decision analysis leads to neurotic
behaviour; refusal to use decision analysis every now and then may be
reflective of fanaticism (or plain stupidity).

5. Answers to a questionnaire on the training of decision makers

The questionnaire, whose items are repeated here, had been sent to
the participants prior to the conference. The answers provided here
logically follow from the line of argument held in the preceding part
of this paper, and may therefore sometimes sound repetitious. In
fact, having arrived at this point, the serious reader should be able
to answer the questionnaire items himself, and come up with very
similar results. (The left hand reference (a) (b) ... is the ordering
of the final questionnaire, namely a to h in the Introduction).

```
Introduction
Reference
```

(c) Q1. Does it make sense to think of decision making as a skill, or a
collection of skills?

A1. Decision making is based on a collection of skills, whose size and
contents heavily depend upon the type of decision problem.

Introduction
Reference

(e) **Q2.** What is it that the decision maker needs to be taught? Concepts, facts, principles, attitudes, procedures, heuristics?

A2. The decision maker seems to need training in :

(i) Rules for the formal classification of decision problems;

(ii) How to formulate a theory of the task environment;

(iii) Rational procedures for pre-decisional structuring and evaluation;

(iv) How to identify consequences of prior actions;

(v) Basic psychological characteristics of himself as a decision maker;

(d) **Q3.** Can decision making skill or skills be developed in a sufficiently general way that they can be applied in a variety of specific contexts? (A limited set of contexts?)

A3. To a large extent this would seem possible, given that decision makers have been trained sufficiently well (cf. A2). But decision making efficiency will depend upon the decision maker's familiarity with the relevant task environment. This seems to be the greatest restriction to generality.

(f) **Q4.** Can decision making skills be developed through a 'learning to learn' type of exposure?

A4. Teaching the distinction among stages in a psychological decision process as well as a proper classification of decision problems may provide fruitful conditions for learning how to do it yourself.

(b) **Q5.** If individuals can be trained, how can that training be accomplished most effectively?

A5. By means of courses covering the points mentioned under A2; by way of computer simulations of simple and complex task environments; by means of 'apprenticeships', during which they have to carry out numerous decision analyses with genuine implementation of the selected course of action (a risky game!) and continuous monitoring and evaluation of actual consequences.

(a) **Q6.** Can individuals be trained to be effective decision makers in unstructured (and structured) situations?

A6. Yes, see above. But this training will be most difficult for unstructured situations. The latter will be structured under the influence of the decision maker's subjective theory of the task environment; selective processes in perception and memory and biases of cognitive simulation may play a larger role here than when the situation has already been structured (due to simplicity? by someone else? out of custom?)

> Introduction
> Reference

(h) <u>Q7.</u> A corollary approach to both training and decision-aiding is
job/task design. How should decision tasks be structured?
<u>A7.</u> Whoever is going to structure them should be an unbiased and
unlimited decision maker himself. The best solution is to let
decision makers structure problems themselves according to a formal
classification of decision problem variables, so that they know which
methods and techniques can be used, which special difficulties may be
encountered, and which psychological limitations and biases are
likely to operate.

References

Fleishman,E.A. Toward a taxonomy of human performance.<u>American Psychologist</u>,
1975,30,1127-1149.

Howard,R.A.,Matheson,J.E.,& Miller,K.E. <u>Readings in decision analysis.</u>
Menlo Park (calif.):Stanford Research Institute,1976.

Montgomery,H. & Svenson,O. On decision rules and information processing
strategies for choices among multiattribute alternatives.Göteborg
Psychological Reports,1976,Vol.6,No.3.University of Gothenburg,Sweden.

Payne,J.W. Task complexity and contingent processing in decision making:
an information search and protocol analysis.<u>Organizational behavior
and human performance</u>,1976,in press.

Rabbitt,P.M.A. Review of E.C.Poulton's <u>Tracking skill and manual control</u>
(London:Academic Press,1974).<u>Quarterly journal of experimental
psychology</u>,1976,28,506-507.

Toda,M. The decision process:a perspective.<u>International journal of general
systems</u>,1976,3,79-88.

Johnson : One point about training. In many cases it's not important
to the decision maker in terms of future consequences since, for
example, in the military, three years is a normal tour and he's gone.
So that what happens after he's gone is not his responsibility. In
business it's often true also a successful manager moves on, so if
the consequences happen after his time in that position, they're not
important to him.

Vlek : O.K. That would suggest the importance of looking at time
preferences on consequences, and whether and to what extent the utility
of remote consequences is discounted. I have the feeling that something
similar occurs in risky situations, where extreme remote but very
damaging consequences, low probability, high damage events are
discounted.

Goldsmith : The judge in court also receives only very limited feed-
back in correcting his decisions.

Kaplan : The consequences of his decisions may have serious effects
on the organisation he leaves behind. Would you now consider that as a
critical point despite what you say?

Johnson : Maybe he makes good consequences whilst he's there and then
bad consequences for later which he can discount.

Vlek : Take political decisions which are done on a four year basis;
four year time span.

Pask : Speaking as a psychologist-cum-system theorist rather than as
a decision theorist, per se, I find your approach extremely attractive
because it is obvious, I would have thought, that decision making, what-
ever else it may be, does involve the whole of the psyche and the whole
of the immediate and accessible environment.

Vlek : Not the whole. Maybe a good analogy here is to distinguish between the hardware and software of the human cognitive system. The hardware being the basic characteristics that operate in any cognitive task, the software being the program which runs at a particular moment, and when we could say that the hardware is generally applicable in all situations, in all tasks, but the software should, of course, be tuned towards the particular decision task at hand and it will be. In any process you can recognise the basic element of the cognitive system operating even though you can only make rather gross predictions about task performance.

But at the same time this cognitive map may be, should be, a dynamic cognitive map.

To what extent should we try to develop tasks specific theories of behaviour, rather than formulating general theories which we can apply to any situation.

Johnson : You're raising a different question when you say you have task specific theories.

Vlek : Let me elaborate that a little bit. There have been other people with whom I had a discussion recently about the non applicability of theoretical notions that came out of research on so-called higher mental processes, notably problem solving research. It is one of the conclusions of Newell and Simon's work on human problem solving that they were surprised to find that 95% of all the observed behavioural expressions were task dependent. Only a very limited amount of behaviour was invariant across tasks; that we should stop trying to develop general theories of cognitive behaviour and should try to become task specific instead. That seems very much like your point of view and I think I have to disagree. But the usefulness of this suggestion is that it brings us to a sensible deployment of our research efforts. We should not give up efforts to formulate general theories about behaviour, or to get at general characteristics of any cognitive system. But at the same time we should try to get out of the laboratory, because we should become much more conscious of task dependence, task specificity, in the use and the combination of these general characteristics.

Zeidner : If I recall correctly from the paper, I'd say that you became disillusioned or disenchanted in running some of the isolated laboratory experiments; that you stepped back to try and understand conceptually where you were going. You are at a point, now, where you could have embarked upon a series of specifications which give you a greater strength of purpose: a paradigm in which you have more confidence yourself.

Vlek : I think that this set of ideas that I've worked out for myself contains a lot of mandates for specific experimental work. But I'm a little bit afraid to get back into the very simple paradigms which have a lot degree of representativeness.

I want to illustrate with the third division of the taxonomy of task variables, the availability of choice alternatives. In the tradition of decision making research, this variable has been greatly neglected, although there are some normative studies, and it seems to be a very frequently occurring situation that choice alternatives come to a decision maker one after the other, that a limited set consisting of more than one are retained. Due to memory limitations alternatives are forgotten but others appear; in any case the whole set is not available all the time. These alternatives may be of any degree of complexity. They may be dynamic in character and have a multi-attributable utility structure. I also think that we need a totally different, untraditional decision criterion here, because the decision maker has to carry something like a satisfying criterion against which he checks any new alternative that comes up. The question that any experimenter asks of himself is, 'how can I start to look into this problem'?

The simplest situation to investigate is to present subjects with numbers, representing amounts of money, and to say, well, I present you amounts of money one after the other, and you pick one and be sure that you pick one before the 10th, because the 10th is the last one. All sorts of variations can be imposed. Such experiments produce results which will make other people curious and these experiments are a great relief for graduate students looking for a dissertion project; but overall we have another "book bag and pokerchip" type of experiment, where subjects are being confronted with a highly simple situation which is relevant only for its sequential character; as long as we don't take all the other aspects into account, we were stuck with that simple paradigm. That is what I am afraid of.

SECTION 3

3. <u>GAMING AND SIMULATION, GAMES AND TRAINING SYSTEMS</u>

<u>Overview</u>

Two papers are classified under the heading. One of them, presented by Dr Kidd, describes management games of varying complexity that the author has used for training. The other paper, by Dr Broadbent, is a detailed and personalised account of observations made whilst games of this type (in one case, an identical game) are being played.

Both papers deal with cognition and behaviour related to decision; in that respect the present category of material is artificially distinguished from "Cognitive operations involved in decision taking". However, the two papers hang together, since one sets the background for the other.

Where topics overlap, the findings back up the general conclusions of "Cognitive operations involved in decision taking", and particularly in reference to learning. For example, there is an agreement that the value of games as training aids for decision is due, in a fairly straightforward way, to task familiarisation but that learning to decide (in contrast to learning a response mode) depends, when it occurs, upon something more than simple training expedients like knowledge of results Subjects should learn the techniques of gaming over a difficulty graded series of games. Remedy of outstanding deficiencies depends upon prescriptive feedback or, at least, diagnostic discussion with the participants.

General findings apart, so much of the content of this section comes from personal and closely monitored observation of what really goes on, so much was also elicited in discussion, that it would be difficult for anyone except the observers to give proper emphasis to the rich but diverse comments upon learning (and not learning) that are reported.

"The Study of Complex (But Not Yet Real) Situations by Management Games"

Kidd : My task this morning is to try to proceed through the abstract
that I gave to you all. I will tend to talk mainly from these slides
(the figures) so let me introduce a task that I present to a few
players at University. It is a mini business game in the sense that
we only have six decision variables to play (Fig.3.1(a) and Fig.3.2). The
six variables (as you see in black) to buy some plant, to schedule it with

normal time working, to schedule with overtime working, to hire men, that
is salesmen, to set a selling price for the single item and to set an
advertising budget to push the item in the market. There is a single
market and it is a very simple response system. The time base is as
shown. The asterisks are the decision points set against the time base,
and decision points like that pair (plant and N/T schedule) will give
work-in-progress and inevitably stock, which is sold to yield sales, a
debtors list and finally, money comes in. This is over a total time
period of i to i + 4 time periods at quarters. Which in terms of the
game initialisation, is time period one through to the fifth time period.

The only data I give to the players during the week prior to running this
exercise is that they do have £150,000 in cash and a diagram (Fig. 3.1(a)).
I ask them to write a brief statement, 150/200 words, upon their plan of
action, and they are allowed to think about that for a week. They then
come along to a computer terminal and play this game for about $1\frac{1}{2}$ hours
to 2 hours, depending upon their total decision time. They play through
from quarter one, spending their cash in hand, to quarter eight; and
between each quarter of play there is a full set of financial output.
At quarter eight I stop them so that they can go away for another week
to consider the data they have self generated, and after that time I
expect them to hand in a protocol to verbalize how they were trying to
optimise. I ask them to state what the optimal series of plays will be
in the next four quarters. That's the third year of play. I input
this data myself to stop students fiddling the data on a realisation of
what might be better. I review the outputs, both in terms of ordinary
financial decisions and Lens analysis, which we will be mentioning later.

This game was also used as a training session for a much larger game.
The larger game in fact, has the same structure as the existing mini

game (Fig.3.1(b)).Now the players have raw materials to consider, and there are research and development investments which may yield a product breakthrough from which they get enhanced sales. There is some financial management such that one's debtors can be brought forward and create a cash input one quarter early, also bank loans are now negotiable in any quarter. So one must maintain much more of a business atmosphere in this larger game. Whilst it has still a single product range, there are now 19 markets; 19 territories which creates the greater complexity of this larger game. Thus the structure is the same but the marketing complexity has increased tremendously.

Just go back to that diagram(Fig.3.1(a)). What would you do as a student? What sort of strategy would you use given only their information? And having perhaps the knowledge that I, as a teacher, might be slightly perverse with you as students. Any ideas?

Pask : It's the hidden perversity which makes it so difficult.

Kidd : Yes, well I think that is incorporated within taught courses. Now to return to my original questions: are there any natural inclinations of yours, if you were sitting at a terminal, of your method of sorting out input decisions? Or, would you just sit and type as a query comes off the terminal for data entry?

Phillips : How long would I get to do this?

Kidd : One week. And, of course, although I offer one week, I am sure it is only a few hours for students, given their usual mode of work!

Phillips : Do I get to the game: do I get to muck about with it for 'free' first?

Kidd : No, you should think of it in a straightforward, dry, cognitive manner.

Phillips : No, I don't. I'm very suboptimal, that's why I study decision making, I like to muck around with it for a while until I get a feel of what is going on and then think about it.

Kidd : Yes. Well, that's the sort of answer I want. But students' eventually have to muck around for real. What I'm looking for is designed mucking about in this same situation.

What I'm trying to do is to supplement my teaching in Operation Research and Systems in the fourth year of the students' course. They, the students, ought to have some sort of ability to collect their skills together to be able to play this game. I'm using it as an integrating device. <u>They</u> ought to be able to express a problem formulation to use statistics. Instead of somebody else directing them to analyse data. They can, in fact, create data for their own explicit analyses. I'm looking for that sort of awareness: what sort of data to collect. Similarly for the accounting data. It is unlike the lectures in accounting: they now see data derived from their own decisions, which is back to my stress of gaming once more. They take a decision and an answer arrives. But they've also got to undertake a realisation or projection into the future.

As for the Research side, it is pretty well coincident with the teaching aims. I'm looking for students creating systemic concepts, and so on. But in the end, very end, I feel I must run a very good debriefing, so that all the students can find out what they ought to have been doing against <u>my</u> normative model.

Vlek : Could you say a little more about decisions to be taken, how the computer reacts?

Kidd :   First, what size of industrial combine do we want to create? Maybe like this hotel, do we want one this size, bigger one, twice as big. What magnitude of operations?

        Second — schedule, logically. I presume we cannot over-schedule the plant. But you can schedule from zero up to that maximum. Third — overtime, the extra capacity, the overscheduling aspect, like working harder for, say, two more hours per day. So that is a decision to increase the standard rate of work.

First, subjects are trained for one quarter. So all these decisions
have got to be considered somehow, but, of course, I don't say that to
the students. That's my background presentation to you.

The management of the game system is really a series of computer programs
automatically linked together, see Fig.3.3(a), so that "I", really the
computer, will effectively start with a small program "Guide" which
says welcome to this playing system. In other words I am conversational
with the players, to put them at ease. Sometimes they may have not sat
at a computer terminal before. There are two programs "Play" and
"Calculate", simply accepting data, the decision data, and which
calculate profits, then prints cash flows, profit and losses, and
balances automatically at each quarter end. But at quarter four and
quarter eight, ie. year end stages in the first two years of the operation
of this mini game, the system jumps to another program and there I have
attempted some outcome feedback experiments. I create more information
for some players, ie. greater information than the automatic cash
management output. One part of the experiment has no extra feedback.
Another, a data table and a graph. Let's explain a little more about
what these are. Effectively it relates to marketing management. But
the data table and the graph are, I suppose, hints from my side (as
the experimenter) to the players, ..."This is the sort of analysis that
might help you". And for the third level of experiment a data table,
graph, plus an explicit regression equation, which links the sales per
man to the advertising per man and to the selling price. So an equation
is given out, but with no statistical evidence as to its validity.
It is an equation, a suggestion.

Zeidner : Are these feedback programs specific to the enquiry of each
players of the routinised game, or are they more specifically response
sensitive?

Kidd : The output is routine, given the current date is Quarter 4 or
8, but the given extra data is dependent upon a random design, such
that one or another player will receive more or less data. Now the
equation that may be output is dependent upon the players' data input.
There is a criterion equation governing the actual sales given input

decisions upon price and advertising, but there is no attempt to force the players to vary their decision set. So the output equation, derived from the players data, is a more or less good match to the criterion equation.

Zeidner : So this is an enormous undertaking on your part in terms of the programming software, to get this thing going?

Kidd : Well, I think that of the left hand side of the figure 3.3(a), but for the other side, I just cannibalised some existing software for a multiple regression program. It's fairly simple. Therefore, as a program existed, I just put inputs into it and got outputs out. And there was slightly different file handling, etc.

When this series of programs is used for training, prior to a larger game, further programs were automatically chained together, so that from a regression program we would get data table, graph, and an equation, see Fig 3.3(b). The further offer of stepwise optimisation and the offer of financial simulation, which wouldn't update data files, but is based upon historical data. Players could, as a group, quite easily choose not to use these extra programs and go straight back up to the "Play" program, saying, "we don't want to know anything about that!" But most players of an enquiring mind would try to look at some possible futures through optimisation and financial simulation modelling, and then back to "playing" their pseudo reality. This sequence, when used as a training for the larger game, took about 2 hours, with 4 or 5 people sitting around a terminal. The smaller game, when I use it in an experimental way, as first described, the first 2 years takes approximately $1\frac{1}{2}$ hours per person.

So that is effectively the mechanism of what happens. Within the game structure, as players proceed through time they learn. They learn about the range, magnitudes, costs of their various decisions. And they complain to me in the experimental situation that we ought to have had this data before. Well, naturally. But I don't see why they shouldn't, in fact, create an algibraic model of the structure; typical of text books costs are C1, C2, C3, decisions are A1, B1, etc. I don't see why they cannot do that, but I find, and I agree with them, it is very difficult.

Nevertheless, when they proceed through their playing space, they will find that they can buy plant in blocks of 50 units of capacity. They buy one, two, three, four, blocks, depending upon their own feelings about whether they want to be a big business, small business, and so on. They can schedule virtually continuously by single units of this capacity, with overtime, up to 40% of this maximum capacity. Hiring of men, integers' obviously, but some people try to hire fractions. Well, they say, lets in one point, seven people. Players are upset when they haven't got the point seven. But I automatically hire the integer part of the input - it seems reasonable. Advertising is constrained, you will see why later, up to a total of £1,500 per man hired. Selling price is a continuous variable, I'll accept any fractions, but it seems pointless to me to incorporate decimal parts, that's up to the individual player. I advise, or the output from the program advises, a maximum selling price of £650 per unit. Zero is a natural minimum, but in terms of the sensitivity of the parameters inside the program, even with full advertising per man, a price of anything greater than £500 is unlikely to sell any items. Nevertheless, it's quite interesting to observe the maxima used by players. After printing the maximum is £650, some people say we must input £650, and they are rather upset by finding no sales. So you have to be extremely careful about that information you are feeding to the players. And so I really felt that, something high, O.K., but not £1,000, because some idiot would put £1,000. Some players try anything. So there were the input and output data they acquire by playing. Again, Charles, it's still very difficult, isn't it?

Vlek : You must have a model of the dynamics of this situation to get the computer to run at all.

Kidd : Yes, but I find that is fairly simple, because, in fact, I've inbuilt a criterion equation which relates in advertising and price and gives an output in terms of units sold.

Vlek : And then you suppose that the player is able to formulate something which approximates the model that is underlying your program?

Kidd : Right.

<u>Vlek</u> : And thats, well I begin to find it tricky because you make a
hell of a lot of assumptions about what people already bring with them
when they come to that situation.

<u>Kidd</u> : Yes, I presumed in the experimental setting that the course has
sufficient structure underneath it to bring us, at least in the academic
sense, certain skills.

<u>Vlek</u> : Which means that your experiment would probably collapse if
you would use the usual first year psychology students.

<u>Kidd</u> : Yes, probably collapse, though this is not certain. The variety
of decision inputs quite often creates an unlooked for experimental
design. So that one can still or often regress the marketing data. It's
not a <u>design</u>, it just happens. But what I'm really looking for is some
analytic statement of this situation. Players could be perverse and try
to just act out against the experimenter, but I think that there is some
sort of underlying human system, which, if not going for profit maximisation,
somehow tries to make a profit bringing cash, so to survive as an ongoing
concern. They are very willing to 'play' with the experimenter insofar as
they themselves get taken over by the game, rather like playing the game
of Monopoly. In fact, games were looked at in an article a year last
February * where there was vast, not regression but correlation analysis
leading to path analysis undertaken on all sorts, all types, of decisions
inside a business game. All inter-correlations were computed, finding
out what really was the pointer system and profit was it. People were
bashing away at whatever, but, the answer was profit maximisation or at
least, attempted maximisation.

* Hand, H.H. and Sims, H.P. "Statistical Evaluation of Complex Gaming
Performance" <u>Management Science</u> 21(6) pp 703-717

Whatever players said they were doing, profit making is the answer. Therefore I would think, why not just go back one stage, cash asset management really is this lead into the profit maximisation. And in this simple experimental system, what is it that generates the cash? It's simply sales of the one commodity. We are selling at a price, people pay and cash comes back into the firm. So we generate an <u>expected</u> sales volume. So you push the items onto the market, people buy them, and this gives you cash. You hope it's the right balance of selling price etc. for profitability. You hope, obviously, but if I give no data to the players, then their management in the initial plays of this is either fortuitously good, or terribly wrong.

This data is all I'm interested in at the moment. They're generating data, and why? Well, the game itself and many organisations split into two reasonable parts. You're making the product and you're thinking about how to get rid of it, and a fairly simple balance occurs. Don't sell it too hard, otherwise you run out of stock, because production hasn't made it. Don't overproduce, otherwise you have a great big stock and that costs money as well; hence the balance. To remind you of the decisions they're taking – the 6 decisions I mentioned earlier, men allocation (or hiring, really), selling price, advertising, plant size, scheduling, raw materials. The only way to link marketing and production is by some sort of analysis. A magic word, yes, analysis. I don't (really) care what sort of analysis the students do. It could be just looking at their financial systems. Or working out, when they made 'that' decision, what it cost, and what the result was. But I would think that it's much more powerful to look at the marketing side in terms of an analysis of the expected sales, and try to do something with an equation of expected sales. The argument stems from a review of the time base of these decisions. We go back again to that first diagram (Fig.3.1(a)). The marketing decisions, price and advertising, yield a result sales straight away. Just over one time period. Over-time, the increase in production yields some sort of result with respect to potential cash input, with one extra quarter delay. Plan, normal time scheduling, hiring of men, yields results in terms of cash, or potential cash, input over a period of three quarters. So the sensitivity to changing your potential future is reliant in the short term on pricing

and advertising knowledge. Change the pricing and advertising, then you change your potential debtors even though eventually the cash will come in. I'm not too bothered in this training game about the potential to get into an overdraft situation, ie. to go negative with respect to cash balance. I know that is a bad situation but I don't want to have another decision (ie. bank loans) to think about. So all I'm looking for is that the firm wants cash, we'd better gain it with those available variables.

When players sit at their terminal most of them will use this diagram (Fig. 3.1(a), as a decision aiding diagram. And this is wrong. And I think it was Charles who mentioned temporal confusion yesterday. Its very easy to look at this diagram and say, fine, we take these decisions and it comes forward to here in terms of a time structure, quarter one of the year we take these decisions. Quarter two, we can take these decisions, built on quarter one. Quarter three we can take these decisions built on quarters one and two. But at the same time, when we are in quarter two, we can take a _new_ decision on plant, a _new_ decision on normal time, a _new_ decision on men. When we are in quarter three, we can take a _new_ decision on plant, a _new_ decision on normal time, a _new_ decision on overtime from last quarter, a _new_ decision on men, and so on. In other words, by the time you get to quarter three, you can take six decisions which relate forward in terms of plant, normal time and men to something three quarters ahead. You can take a decision on overtime which looks backwards one quarter, you can take decisions on pricing and advertising which looks back three quarters onto the generation of stock. So you've got forwards and backwards to think about in your mind, to juggle, to get an attempted balance. An attempted balance of marketing and production. The players will then confess, sitting at their terminal. They will turn round to me and say, "you're a bastard, aren't you?" Simply because you've (the experimenter) got us into this very silly situation of knowing no good data input, magnitudes, costs and so on. We've got to a position where we have, in our minds, a known inbalance between marketing and production. And players' have to attempt to go forward. As far as I am concerned they attempt to go forward by spurious experimental designs".

__Vlek__ : Well, I think that things can be clarified a bit when you talk
of the number of simultaneous decisions to be taken about different
things with different time perspectives. The thing is, that if you
produce, you should have a more extended time horizon, than when you
take decisions on marketing. Marketing is meant to get rid of things
that you produce in the past. But current sales should change your
probabilities of future, expected, sales and these probabilities are
relevant for the production decisions. I don't see the confusion if
you clearly bring the whole thing as a set of simultaneous decisions
on different things.

__Kidd__ : Right, I accept your argument very well. In fact, its the
argument you have to accept. Unfortunately the players become very
confused, because of the lack of prior analysis and a lack of the
instant of playing the situation, of understanding this system in which
they are situated.

__Vlek__ : But that's because you bring them into a number of parallel
decision processes.

__Kidd__ : Yes, that's right, but consider this diagram (Fig.3.5). If
they create this type of diagram, it removes the temporal effects. I
know its a mixture of cash management and utilities and various other
things, but I feel that this should guide players into a knowledge
that the only way of bringing these two major aspects together is
through marketing data generation. The data generation, if you look
at it before you attempt any analysis work, should have been of the
form of a designed experiment. We ought to have high, low, medium,
advertising, high, low, prices, given those are our two decision
variables. These ought to be arranged so that we can attempt regression
analysis. Why do we want to attempt regression? Well, we can be very
clever with calculus, or stepwise optimising, to find quasi-optimal
selling price derivations which will imply the expected sales volume
at that optimal level and will imply production decisions. So, on the
longer term, based on that data analysis or any hint from the players,
I can find out if they know really what they're trying to do.

Pask : When do they get to the steady state in respect to the situation? Obviously the situation opens up over the first few quarters of play, then a situation of maximal complexity is reached, and appropriate look back and look forward horizons are determined, and you can trade them off trying to do a classical experimental design against loss, because if I was to vary one variable at once or two variables at once, I would probably sub-optimise. It may give a good analysis, but I'd sub-optimise local performance. What I wasn't quite clear about was what their target horizon is, and if it actually reaches a steady state before then, and is it assumed that the firm is going to continue operating indefinitely, or is it going to operate for five years, or what?

Kidd : The target assumptions are left vague. There is the knowledge on the players' part that we will stop interaction at the end of the third year of play, ie. 12 quarters of play.

Pask : The firm doesn't necessarily stop, even though the fact is that you take another job.

Kidd : No, that's right. And there is nothing to be gained or lost by driving the firm into the ground. Normally I would say one is looking for a potential continuation of the firm.

Pask : If survival is the thing. That's perhaps why I was pernickity about optimising profits which is not always optimising survival. It seems to me this would lead to rather different experimental designs.

<u>Kidd</u> : Yes, perhaps it would, but on the other hand, given the simple
response function inside the programs, the marketing experimental design
could be trivial or multi-variate. If, in fact, students often work at
a low level, varying one variable at once, to see what happens. Then
they will run out of experimental time, because they only have quarters
three to eight available to generate data. Note this slight change to
Brunswik's Lens analysis (Fig. 3.6) in that here we have cues, leading to
or leading from criteria on one side, to the task response on the other
side. The game is unlike the usual Lens analysis insofar as I do not
present any cues to the players now. They have to set their own. And
that is a peculiarity of business games. One stimulates the environment,
pushes it, and the appropriate answer comes back. This is the realisa-
tion of a reality from the players' viewpoint. That is why I was rather
interested yesterday in inter-correlation; because, for instance, within
very bad experiments, like full advertising and a full price £650; no
sales! That's fine, so you drop the advertising and you drop the price.
And very strangely you can get a series of stimulus variations that are
in a beautiful ratio with respect to each other, and it's impossible to
run a multiple regression analysis simply because there's a pseudo
variable, or more exactly, the problem of multicolinearity arises.
Well, sales are the usual response out of the program but based on
costs (production costs which they have created for themselves via
their scheduling of N/T and O/T), advertising and selling price in
fact a stimulus but not quite the usual Lens stimulus. Profits are
generated by these sales which is the response really that's required.
Not the sales, this is a trivial intermediate. If one can review the
stimulus cues (the response with respect to sales), one can regress.
The data, if there is a good design there. Assuming the background of
the business game is reasonably stable, like reality, one can set up a
regression equation which links across here (left hand side of Fig. 3.5),
and looking at the futures by making the gross assumption perhaps that
the future is going to behave like the past, one can get to predictable
sales volumes and profits. And if you can get to predictable ones, you
can undertake analyses on that equation to get to quasi-optimal results
by differentiating this equation with respect to profit. And at that
stage, I am quite happy to look back to the criterion. In other words,
if your experiment here was very good, then the subjects somehow have
mapped the criterion space and their performance with respect to the

task side of the Lens, looking for profit maximisation, is, in fact, maximising the criteria values. So I am reviewing the usual Brunswikian errors, the actual to predictable error, actual to optimal, actual to criteria errors. But I am using the total modulus errors and not correlation analysis. Because the data are so few here a correlation is not operable, so I look at a total modulus error, which is crude but quite effective. For instance, the number of sales here might have been 30, the predictable sales 35, difference 5 sales. Whereas the criterion might have been 43.

The reasoning I could expect the students to develop from their initial training, is to make an assumption of the market situation, the sales per man is related to advertising and price. The profit per man is related to the expected sales, multiplied by the income less the costs of generating those sales. And differentiating the profit, usual maximisation one can get a point of inflection in the price with respect to advertising showing provided I have that linear assumption, so that zero advertising or the maximum advertising (the constrained variable) is optimal and dependent upon the cost of production, then I will either have an optimal price given by that equation at zero advertising or at fifteen hundred pounds. In other words, if these data(Fig.3.7) were the developed regression coefficients for that equation, my point of inflection would be £258. If the selling price was set at £250, it would be optimal to have really zero advertising and instead of £250 which might have been a player's response, he should have used £216, as the optimum price. If they wish to fully advertise to generate higher sales volumes, the optimal selling price is now £323. But that is the quasi optimal based simply on the regression equation, and it depends upon how good the experimental design was, nb. there are good or bad derivations of these beta values. And so its perhaps a question of how integrated is the whole of the teaching program behind the students. But on the other hand, its a question of their ability to look at uncertain situations, to think about what might happen up to at least this level of abstraction and say what we need is an experimental design. In other words, I hope they are thinking about the whole of the system problem rather than the playing of decisions one by one. All that I have to

do is really for the last year of play is simply input new advertising
and price decisions that the players have developed from their studies
of the data that they've acquired. As well as production scheduling and
other decisions, I'm interested really in advertising and pricing
decisions. What are the ones that you want to use, please try to
optimise the system, give me the data. I feed it in to the terminal
and look at the Lens errors: actual to predictable, actual to quasi
optimal, actual to criteria. And that's one standard, in fact, of
how well the students have understood the situation in which they find
themselves. Another standard is to question if they describe what
they were trying to do, in terms of some protocol, and whatever was
the protocol used to analyse the system and to generate in the last
year of play their decision points. The answer usually is no.
Infrequently do the students understand that they have to look at the
total system rather than just playing decisions in. They don't
understand the real need for a good experimental design. Quite often
they place constant advertising levels and one of the only reasons for
having regressable advertising is that they forgot that they had a
change in the manpower pattern. And so the advertising per man varies
by error. It may be divided by 20 men instead of 10 men, for example.
Thus, they didn't design consciously, it happened. So that creates
data variation.

Pask : How much data logging do they have when playing at this terminal,
they presumably make notes, draw diagrams?

Kidd : The students? They keep it in their minds, but if they
wish to make notes they can stop the terminal sequence for a while.

Pask : Do the notes, sketches, etc reveal much about the decision
process ?

Kidd : Some students do, in fact, create a series, a design series.
Perhaps it is a very bad design.

Goldsmith : Do they have sufficient time that they can do that without
disturbing the progress of the game?

Kidd : Yes, but of course, as you know, if you're sitting with a
terminal you feel that the computer is driving you, you feel you must
create an input, although the machine waits for them. On the system I
use there is a facility for a timed input. So that if in 10, 20, 30,
up to 200 seconds, there is no input, it could automatically provide
an input. And record for me the decision time. But I don't use that,
I just allow a free input, free time. But students do feel driven.
That's all very well I think, in terms of a small experiment.

When I use this game in its slightly enlarged version to act as a
training program of the larger business game, I set the criterion
variables, the internal variables in the program, so that the optimum
selling price was £300 and the advertising maximised at £1,500 for the
best profit. And by the stepwise optimising program the players, as a
group, rather than individual, would be able to generate a neat table
of price versus advertising showing the response surface of profit.
So this is a bit of learning. They have obtained brief lectures at
the management school about regression and about accounting, and so on,

so that people with varying skills would at least have a cursory
knowledge of other skills. And in this school they were grouped into
playing groups to play their main business game. My role there after
was to act as a consultant for those groups if they wished to use me
to help with data analysis, or just to act as a consultant. Somebody
technical.

I found various interesting points: there was very strong support of
the anchoring and adjustment effect based on the "fact" of my optimal
£300 selling price in the training program. If I did not interact at
all with any of the teams in the main game, then over their restricted
series of plays (they only play for five quarters on the large game),
they would initially set the selling price at about £300, £250 to £300,
then over the series of plays against competition they would reduce
that selling price. I feel it is fairly natural to reduce the selling
price against your competitors. So here was typical transference of
training. If I reacted with the groups to undertake regression analysis
and profit optimisation, then I would suggest, via optimising programs
to the advised groups, that the selling price should go up to about £360
or £380, depending on the main game parameters. So some teams would be
raising their selling price against other unaided teams who maintain a
price or perhaps reducing. So there is a very clear indication that in
a free system they anchor on prices which are not optimal but, if guided,
move to quasi-optimal performances.

The other point is: if I was asked to be a consultant, my advice on
profit maximisation was _always_ accepted. No hedging, complete acceptance,
and perhaps there is an answer to that as I am "a consultant", therefore,
somehow good, I was also attached to the college, therefore, I must be
correct and my analyses were printed by a computer terminal and it must,
therefore, be correct. I took care in arranging the output format so that
it was readable, understandable. I hoped understandable, but I'm quite
sure a lot of people just did not understand any of the concepts of the
multiple regression and the research program, stepping towards the quasi-
optimal.

I know some of the people just accepted data because they felt it was
fashionable. Quite often this was the case with multiple regression.
We want multiple regression they say. So give them multiple regression.

If they wanted it for free that is. And with no explanation. We
know that there are comprehension problems with a multiple regression
output and it's not easily understandable. Teams learned they have to
buy understanding.

A staggering fact over the series of seven games, I aided five out of
eight, five of eight, three out of eight, nought out of eight companies,
and so on (see Fig.3.8).In a retrospective analysis of the type of aid
I gave companies and, assessing my subjective quality of aid, I find,
for instance, here that I would have expected this company to come out
first, it did, and these second and third. Partly this is a time base
effect. The sooner people come to me the more quickly do they get into
this optimising run, the more quickly they acquire assets and as the
adjudication is on a calculated share price (which is roughly we've got
more assets now than we had before), then clearly the faster they come to
optimisation, the more likely they are to win. But this performance,
apart from some very odd rank reversals dotted around, is quite
fantastic.

There are two points, perhaps, one I've just mentioned, aid when requested
was always accepted. Secondly, there are within these teams some very
clever people who are quite capable of undertaking regression analyses
and searching for the optimal values. Sometimes multiple regression was
requested. I hand it to them, but it was only after the game was finished
that these players come up to me and enquired of the optimum values.
Were they £380, £1,500 advertising? Yes. We knew it, they said, but we
couldn't get the players in my group to do anything about it. In other
words, internal advice generated within the group is not accepted, external
is accepted. This apparently is very contrary to Jan Huysman's[*] findings
in gaming. He had an external consultant, really part of the computer
system, driving along, creating advice. Apparently to the other players
that was internal and not, therefore, acceptable, so here in these games
I find I'm omnipotent and rather frighteningly so.

* Huysman, J.H.B.M. (1968) Implementations of Operational Research,

<u>Zeidner</u> : Is that surprising ?

<u>Kidd</u> : Yes, I think so, because there was such a strong adherence to
what I was saying. In the earlier games the parameters inside the main
game were wrong in my opinion because the profit maximisation was up at
£550 selling price. I just drove people up there, and that is a long
way to go from £300. And this was an acceptable but terrific
psychological distance. What I am more interested in really is not
looking back on the total span of people's willingness to come to me,
but why do some companies not come? Here, in S5, I know why, I asked
them to bid for me, and nobody wanted to bid in isolation for my
services. Here (S1, S2) I was free, and I have quite strong advertising
of my presence. Here (S3, S4) I was free with very minor advertising
of my presence, very quiet about it. I just existed and let people
come to me. Here (S5) I say bidding, nobody wanted me. Here (S6, S7)
I cost money and people knew I was around. I presented a statement of
my services. Quarter by quarter it cost so much for this type of
service etc. And my effectiveness in fact, is rather stronger here
(S6, S7) less so free, down to the bidding (S5) when nobody wants me, so
pressure of advertising on my part effects the amount of work that I
create. But even so, in some of these games, even this one (S6), one
company said no, we do not want to see you.

The last bit of data to be presented today is something that I'm really
far less sure about, the whole background of the psychometric scoring
system. Some may know it, some may not, it's a straightforward
psychometric dimensioning of a person's attributes. Sixteen factors
off the Cattell 16PF test (one takes about 40 minutes or so to go
through 200 and odd questions, about how we are). So we get a person's
profile. The 17th factor is from the Watson Glaser Critical Thinking
Test, is slightly different in concept to Cattell's Intelligence factor.
Most of the players sat through these test sequences as another
experiment being run on team building, and so I've obtained this data,
gratuitously. Fortuitously I now know the personality make-up of all
the teams.

And I have looked at the initial plays in the main game, merging together
for the analysis, all the groups I did not aid, the ones I partially

aided (that is from partway through the gaming session), and the ones that I aided fully from the beginning of their main game through to the end. I used the ICL 1900 factor analysis, going to oblique analysis, trying to maximise the hyperplain count of zero, or near zero loaded elements. Generally 6 factor groupings come out and I have related here the first 3 factors of the different groups. And picked up some of Cattell's key words (Fig3.9). One can read all sorts of things into it, but .....

Kleiter : Could you give us some kind of reference on how large are the correlations?

Kidd : I haven't got the data here but, very roughly, 0.5 or 0.6. There's some quite strong correlations taking place.

But now it's a matter of where do I go in this output(Fig3.10) to give some sort of story line here. Taking the first three factors, just looking at some of the cues, and in one brief sentence, the groups I do not aid are a series of business people who have an image of being hard bitten. "We know what we are going to do", they say. Quite assertive. Quite self assured. And they are clever. They can work out the implications as far as they are concerned, between different possible series of plays. On the other hand, the fully aided people again are intelligent and assured. They have quite a lot of individuality but they are creative enough to get together and say, "lets forget about our individual differences, and buy some decision space, lets go outside the group as well". So perhaps the low team integration is somehow containable, even removable, by being quite bright and assured. Finally, the middle lot. They are a bit wary and suspicious of life in general. But they are not quite bright enough (they are at the negative pole of the Cattell Intelligence, which is not to say that they are stupid. It's just that they are less, rather than more, intelligent). They are not quite bright enough to really jump out of their situation, so its like, "well you know, lets wait and see a bit ....."

Vlek : They do the test, when, before, or after, the game?

Kidd : Before. The teams are set up, based on the individual psycho-
metric profiles. I had nothing to do with the team building structure.
I obtained data on team structure after the games are complete.

Broadbent : John, could I clarify exactly what happened here, because
I haven't understood? As I gathered it, you did a separate factor
analysis with a class of people, that is the class who asked for aid.
Then the main factors that came out in that have the labels as you
suggested. All that that shows us is that the variability within that
group is greatest in those factors. In other words, that group includes
some people who are at the opposite end of those descriptions.
O.K. So you're saying that in the aided group, the main differences
within the aided group are of one type and the main differences
within the not aided group are of another type ?

Kidd : For a given set of subjects they all took the questionnaires –
That's about 60 odd subjects in each main game. In fact there are
more subjects as there are replications over two or more series of
plays. All the subjects were grouped into teams under an algorithm
unknown to me , depending upon the test result.

Phillips : Well John, there are two ways to assign. The one is to
look at the profiles to ensure you get _exactly_ the same people in
every team. The other approach is to make _purposely_ different teams.
What was the aim of the selection process?

Kidd : Purposely different teams, under a design over which I have no
control.

Phillips : So, any team is made up of a collection of people who,
sort of on average, have a profile that's different from the next team?

Zeidner : But why aren't you concluding that those teams that
asked for aid are fully characterised by a profile different from the
other two teams?

Kidd : Yes. But I am attempting to extract a new set of factors based on the intercorrelations of the original profile data of groups who behaved in a similar manner, ie. demanded aid.

Zeidner : So the profile has nothing to do with the assignment to the teams? You could have given the profiles at the end of the entire experiment?

Pask : No, I gather somebody else did the assignment into groups that were assigned for some other reason than John's work. But, in fact, the data was available. Afterwards you independently noted their different behaviour: they either asked for full aid or some particular categories of aid. The mixes of teams which have been devised for some other experimental purpose or whatever, were oddly enough, very different in respect of the aid they asked for. Hence one looks to see whether there isn't some kind of trapping effect due to their interaction. You can characterise these people in some sense that makes them a team, and also there are very strong differences between them.

Kidd : Yes, there was a desire by some teams and not others. All the same, I am particularly so surprised by the very strong anchoring effect that occurs in the games. And the need to be very careful in what you are trying to tell players to do. They will take instructions as an order rather than a statement to investigate their reality.

Pask : Isn't this like the undergraduate syndrome; that what is spouted from the lecture podium is a drop from Heaven and written down in a notebook?

Kidd : Not so in the large game situation, where you have iconoclastic and rather hard bitten business men who are quite willing to go their own way.

Zeidner : They are doing a management course and they are back at school.

__Kidd__ : They also carry in transferred skill but it is not necessarily
very strong. They don't take their marketing expertise in. I'm
not sure about transference out, as usual there is no measure upon
long term effects of the training.

__Vlek__ : Well, I just have one more comment to make. The whole
situation is very much in dynamic decision making paradigm for which
we have several models according to the complexity of the situation.
I have been more or less waiting for you to say, or to identify the
decision parameters, basic factors that subjects had to have an idea
about in order to get a notion of the dynamic model underlying the
outputs and the reactions of the outputs to the inputs, so that it
would become possible to identify specific responses or specific kinds
of judgement where things might go wrong or which have to be done with
some reasonable degree of accuracy in order for you to expect a
reasonable optimal output. That relates to my comment at the beginning
that you make such a lot of assumptions about what the subjects bring
with them to the situation. What model they have in mind or are able
to develop about the situation before I think you get to psychological
theorising about the situation. You'd like to be able to lay out the
elements of a model, dynamic model, that they have to go through in
combining all the input information that you display to them in some
way, making this combination at this stage, and another combination at
another stage. In using the output of first stage as an input for
another stage you may come up with reasonable responses but there are,
it seems to me, that there are a whole lot of judgemental problems,
problems of creative designing of the future. Judgements of probability,
statements of expectation which cannot be measured along the way and
checked in and checked out, in order to trace the process that the
subject is going through. This is another instance of confronting
subjects with a rather complex situation. I'm looking at the output
and you see that they don't quite do what your model predicts. Yet
you can't give specific reasons why and where subjects were suboptimal
in their modelling of the situation.

Let me add that I like the game and I think that this is a nice thing
to do with the subjects but in order to get to grips with what actually
happens you should look much more closely into specific stages and try
to do some modelling of the subjects cognitive behaviour.

122

Kidd : I've been trying to understand what, in fact, they are doing.

What I'm really searching for as a teaching paradigm perhaps is to generate interest so people realise that there is some systematic way to nail down the total uncertainty of a system into which they are going. In other words, to experiment with their reality, to generate some data that might help them with their uncertain futures, rather than just acquire data in a haphazard manner, to use a data design.

Vlek : But you don't know what information they need, at which point, and you don't quite know about which relationship they are uncertain and which response is specifically difficult for them to generate. If you knew you would probably have a much more lively exchange of ideas with Berndt Brehmer because that seems to fit exactly into his scheme.

Kidd : There is a very very great difference between the usual probability learning system and this one. There is the need for the subjects to actually stimulate the environment and I am giving no orders as to how they create their own stimulus cue dimension.

Vlek : But you experiment without hypotheses. It's inductive research in order to find out how and you don't want to test whether.

Kidd : Well, O.K. Yet I'm trying to find out what hypotheses the students might have in terms of their protocols.

Vlek : Do you have any ideas?

Kidd : They don't know what the hell they are doing. That's an over simplification. Some people do, some people will say, we assume that we are in a linear growth period of our marketing structure, we're not getting into saturation or overkill. We will assume linearity, we will do such and such.

<u>Vlek</u> : Couldn't this be enriched by carrying out designing and carrying out several little experiments on the side line just to check on, for example, how they relate expected sales to production decisions and just production decisions, so that you know how they behave with regard to that part of the total dynamic system. That sort of thing. If we relate to the discussion we had yesterday afternoon where I expressed my anxiousness about getting involved in another little popular experimental handout, I think that we're at the other extreme. It's very attractive to have such a game but again it's very clear that it produces problems of measurement and tracing the process.

<u>Kidd</u> : What you just said seemed to me not worth while because there are a mixture of very clear deterministic steps the schedule of the play. Only the sales function has any true uncertainty and there is the uncertainty generated as players' attempt to match the beta weights of the criterion equation. The rest of the system is deterministic.

Some of the process is perhaps assumed by the students. The only uncertainty is the marketing side, the cue generation and the derivation of the generation of the output, ie. the expected sales given the marketing stimulation. The only bit that has any uncertainty, the rest is (trivial) addition. It is difficult to relate to your desires Charles, which I accept are reasonable to fellow experimenters.
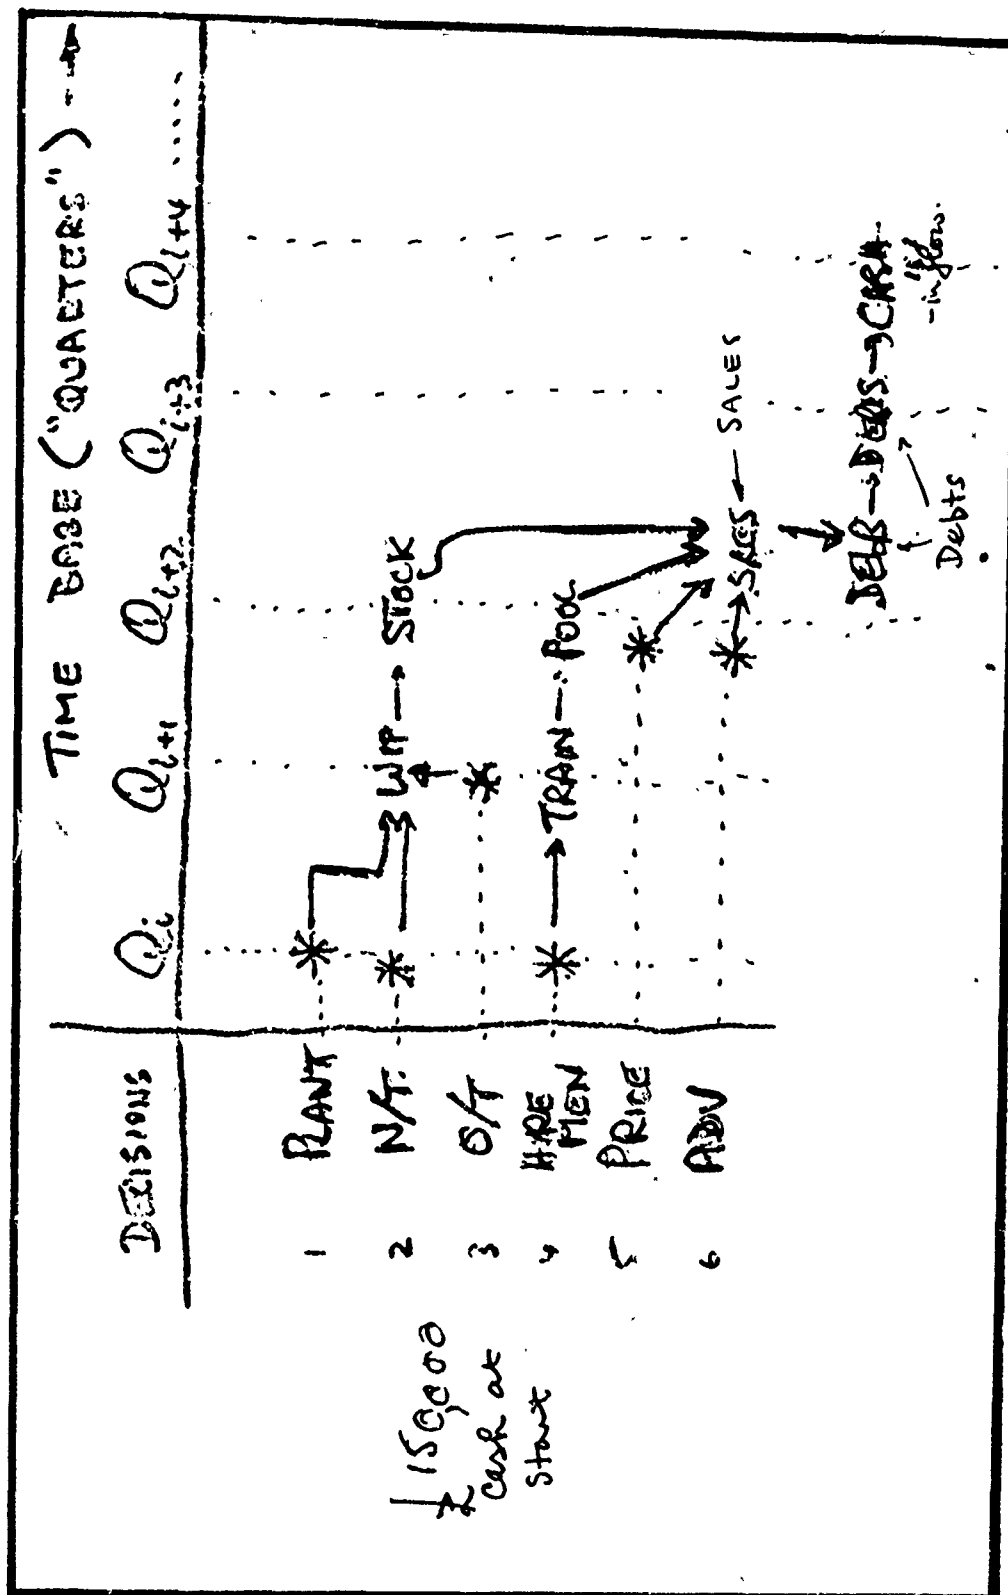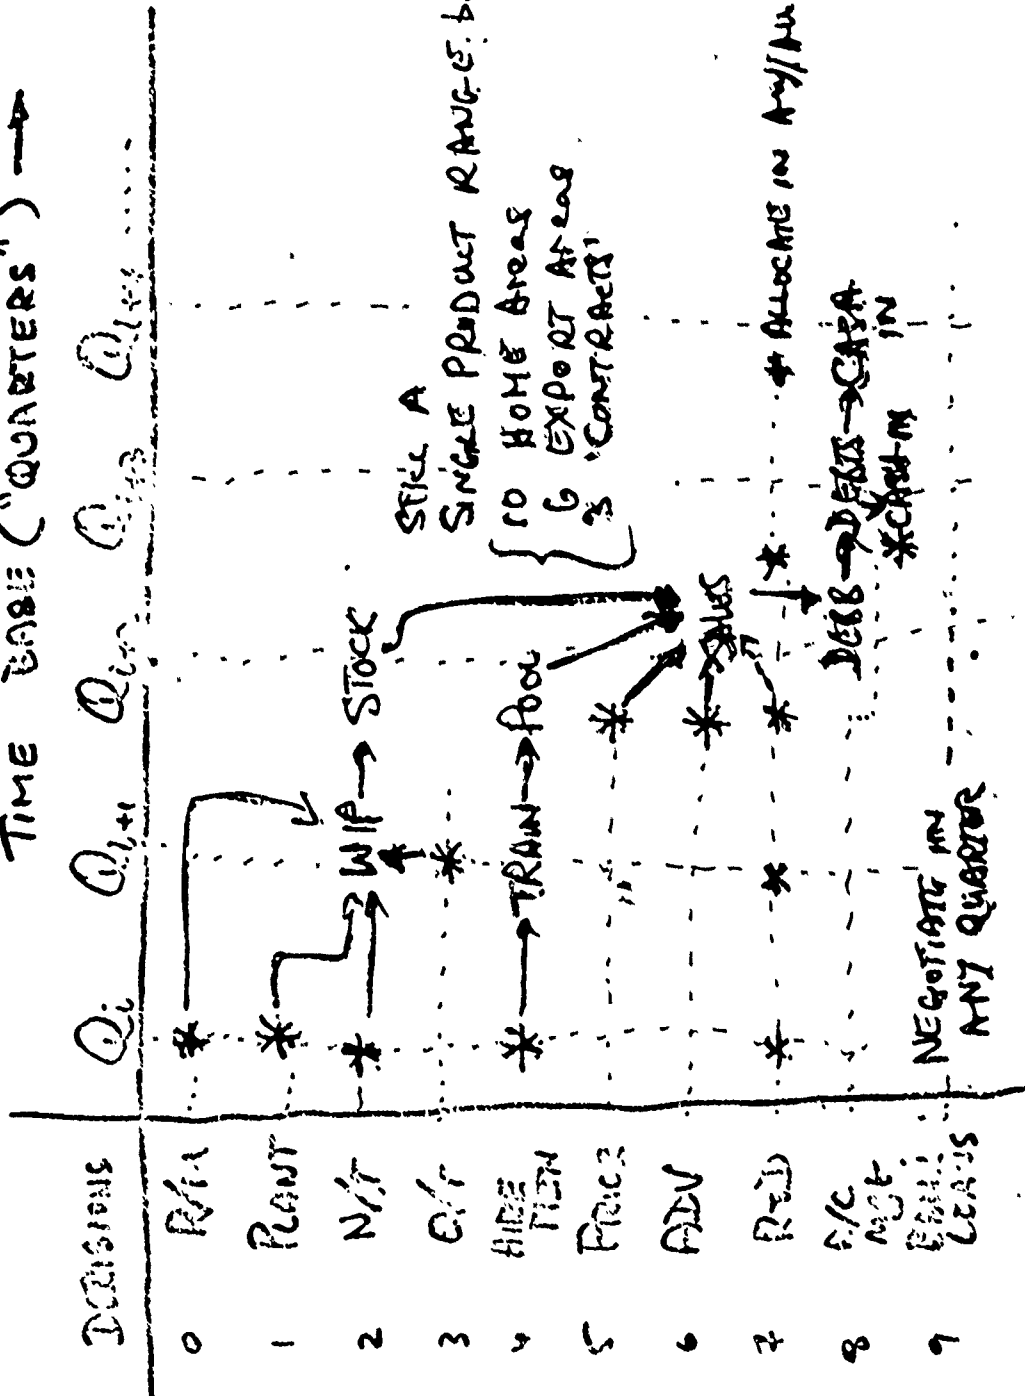
TIME BASE ("QUARTERS") →

Q_i Q_{i+1} Q_{i+2} Q_{i+3} Q_{i+4} .....

DECISIONS

1 PLANT

2 N/T. WIP → STOCK

3 O/T

4 HIRE MEN → TRAN → POOL

5 PRICE → SALES

6 ADV → SALES ← SALES

DEBTS → DEBTS → CASH → ... flow

Debts

£ 150,000 Cash at Start

Fig 3.1a Decision/Results Network for Mini-game.

TIME TABLE ("QUARTERS") →

DECISIONS	Q_i	Q_{i+1}	Q_{i+2}	Q_{i+3}	Q_{i+4} .....
0	R/M				
1	PLANT				
2	W/F	WIP → STOCK			
3	O/T				
4	HIRE FIRE	→ TRAN → POOL			
5	PRICE				
6	ADV				
7	R&D				
8	F/C MGT				
9	BANK LOANS	NEGOTIATE IN ANY QUARTER			

STICK A SINGLE PRODUCT RANGE: but sales possible

{ TO HOME Areas
& EXPORT Areas
3 'CONTRACTS' }

* ALLOCATE IN ANY/IN Q_i

DEBS → DEBT → CASH
STOCK → CASH
→ CASH IN

£150,000

Fig 3.1b Decision/Results Network for maxi-game

DECISION	RANGE
1. PLANT PURCHASE	"1,2"..... (Blocks of 50 capacity)
2. NORMAL TIME	1,2 .... (by unit)... Max Capacity
3. OVERTIME	1,2 .... (by unit)... $\leq 40\%$ Capacity
4. Manpower	1,2..... (integers) ....
5. Total Advertising	0... $\leq £1500$ per Man
6. Selling Price	0..... $\leq £650$ (advised by $\underline{E}$)

Fig 3.2   Decision Ranges in Mini-Game

Fig 3.3a. Linked Programs in Mini-Game

Outcome Feedback Programs.

Management Programs.

REGN

EXPTS:
1) No feedback.
2) Data table + graph
3) $a \cdot 2 + Reg.$ Equation of Sales function

automatic linking when prior game played to with larger game, training

OPTE

SIML

GUIDE

PLAY

CALC

Cash flow
Profit/Loss
Balance of each quarter.

Q1 Q8.
many

onwards to $Q_9 \dots Q_n$.
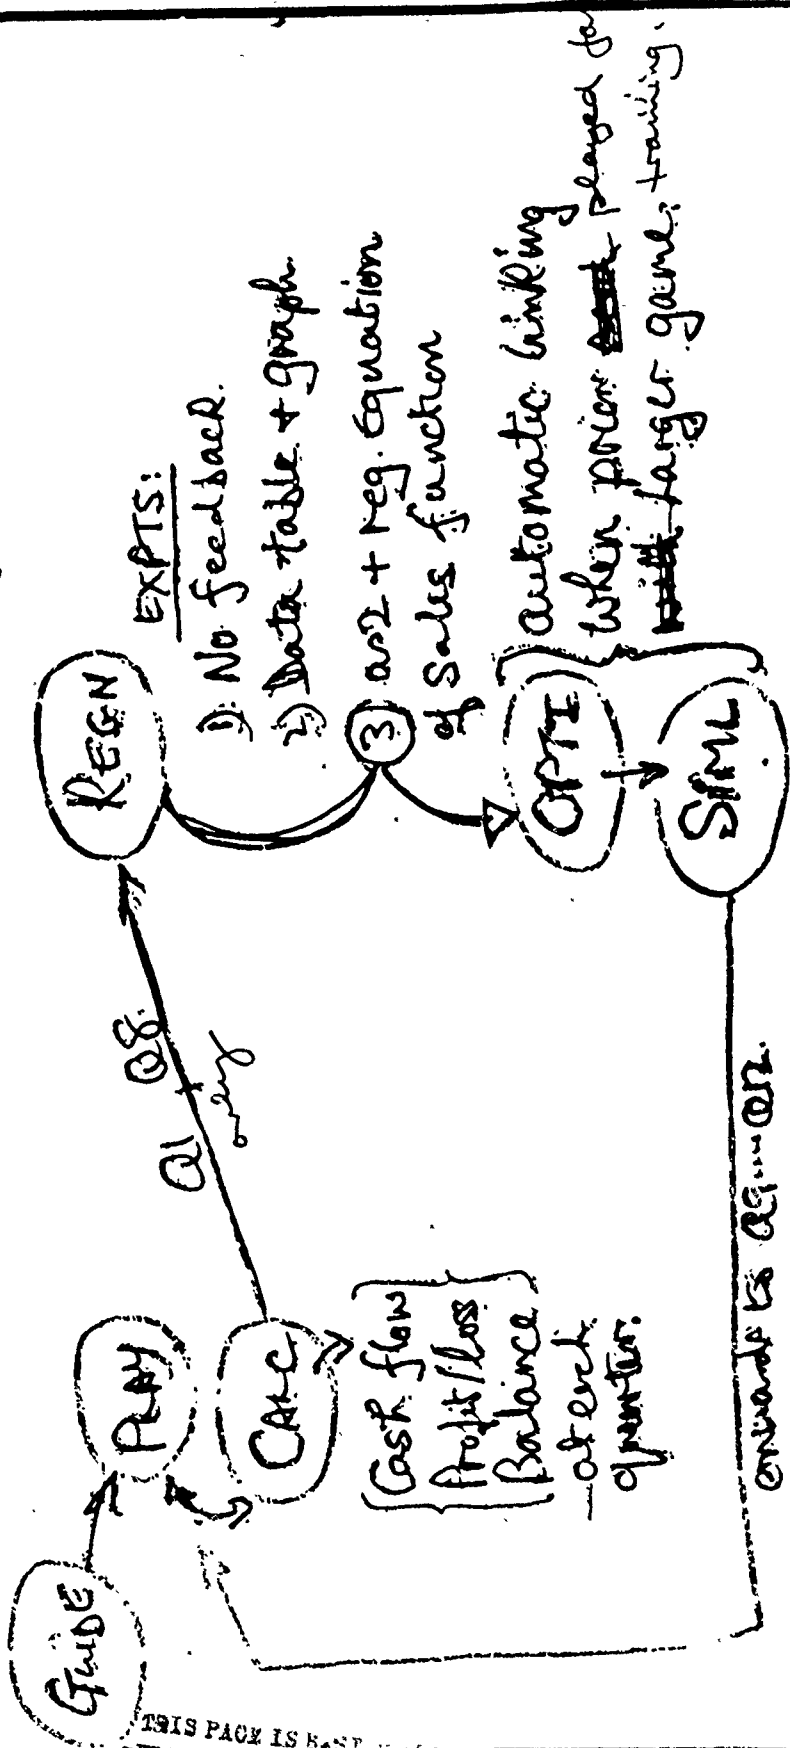
Fig 3.3 b. Linked Programs in Maxi-Game

WHAT ARE MY { OBJECTIVES } ?
{ AIMS, etc }

Teaching aims

Show how to:
- System analysis
- Problem formulation
  (Statistics use (formal)
  (Accounting (formal))

Research aims

Check if players use
- Systemic concepts
- Generate data
- Use statistical arguments
  (Optimisation methods

→ IN ALL CASES RUN A
DETAILED DE-BRIEF to
"Ensure the Pedagogy !"

Fig 3.4 Pedagogic and Research Aims
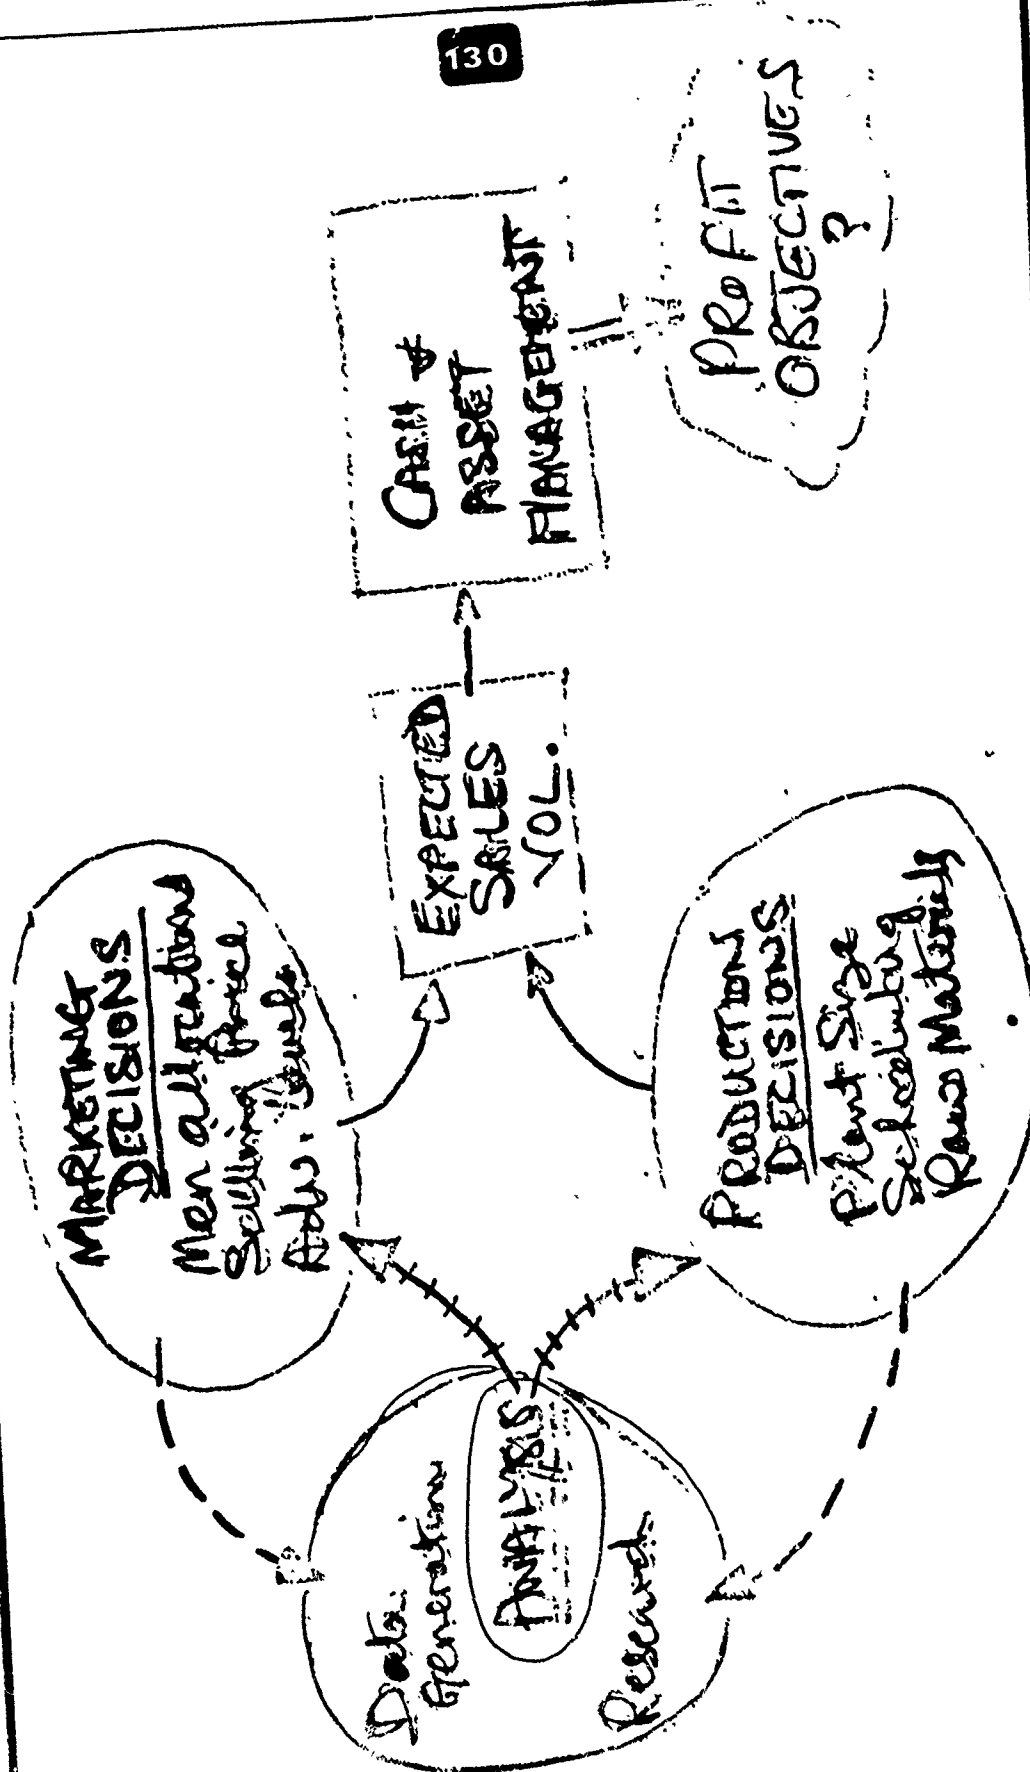
Fig 3.5. Systems Concepts of Mini and Maxi Games
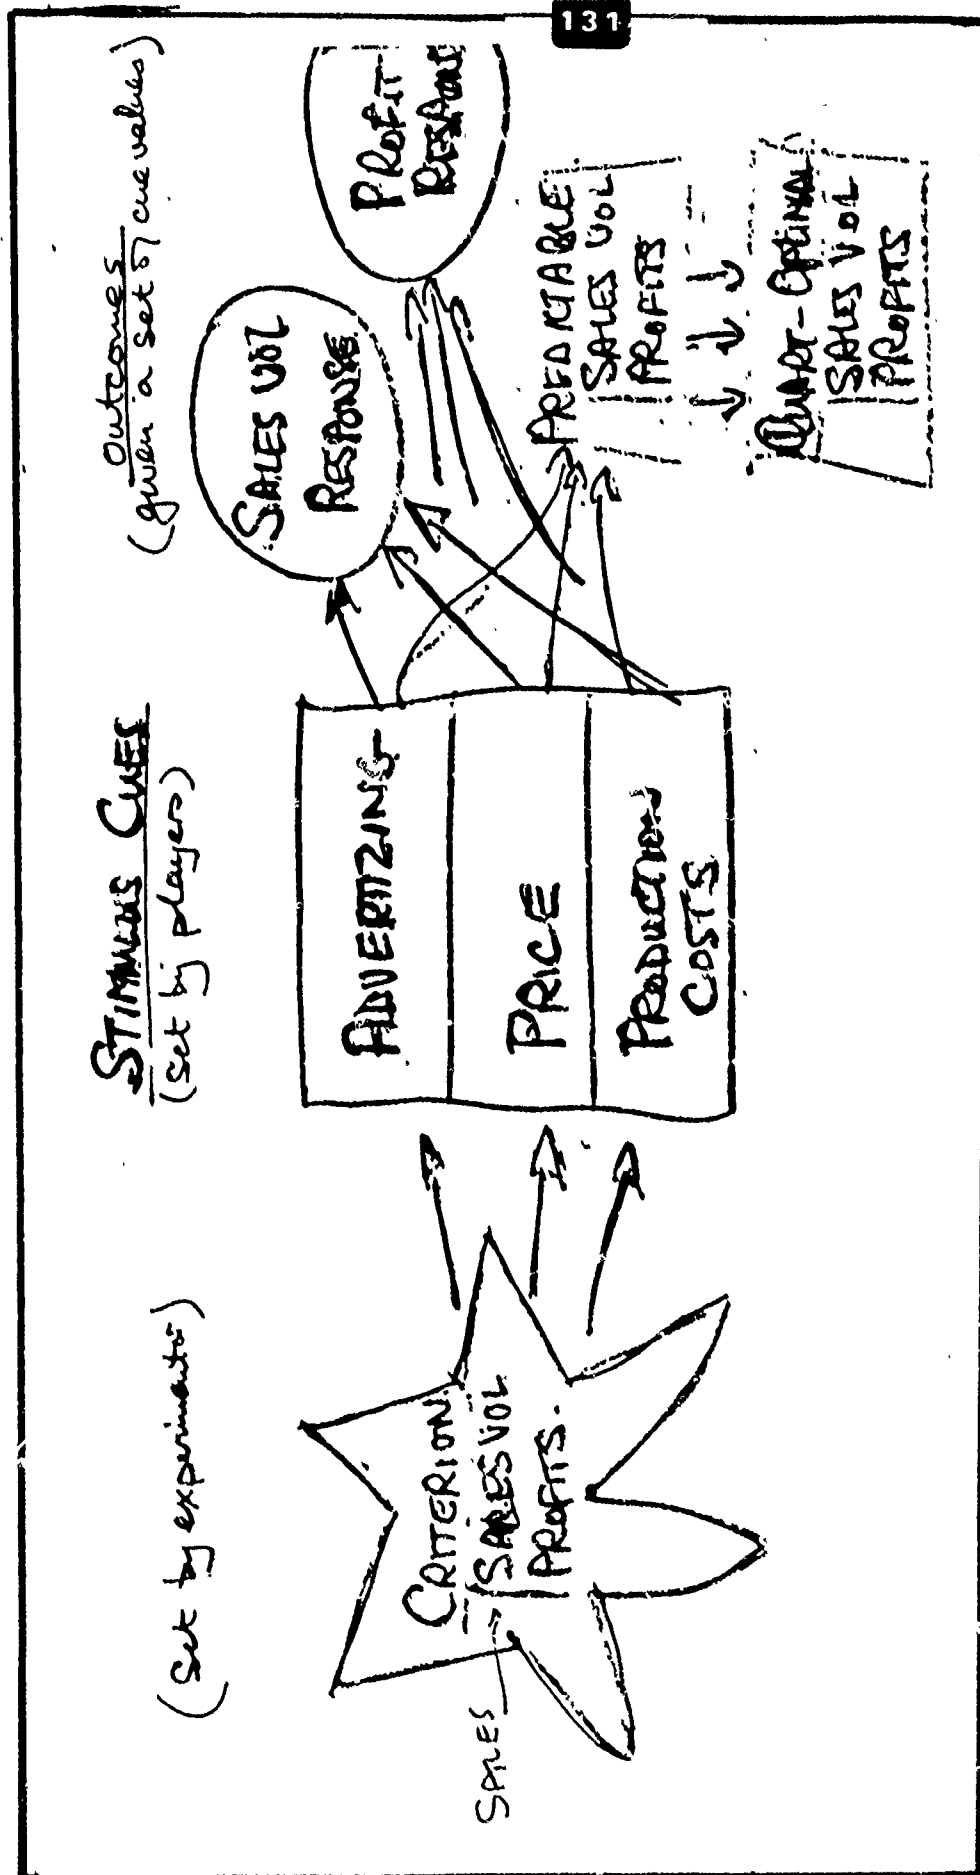
STIMULUS CUES
(set by players)

Outcomes
(given a set of cue values)

SALES VOL
RESPONSE

PROFIT
RESPONSE

ADVERTZING

PRICE

PRODUCTION
COSTS

PROFITABLE
SALES VOL
PROFITS

QUASI-OPTIMAL
SALES VOL
PROFITS

(set by experimenter)

CRITERION
SALES VOL
PROFITS.

SALES

Fig 3.5. Brunswik's Lens Analysis

SALES/MAN = $K_0 + K_1$(Adv/MAN) + K_2(Price) — Assume, and so —

If PROFIT/MAN = Σ(SALES/MAN)(Price-Costs) — (ADV.+ SALARY) Costs per man.

then differentiating etc... to find various important selling prices
accompanying with c = production cost.

$P_I = c + 1/K_1$ ...... the point of inflexion.

$P^* = [-K_0 - K_1(A/M) + K_2(c)]/(2K_2)$ .... the "optimal" price
given $A/M = \phi$ w/1500/man.

If we have :-

with $K_0 = 13.970$
$K_1 = .0056$
$K_2 = -.039$

and $\{ c = £80/item$

$P_I = £258$ — inflexion Price.
$P_g^* = £246$ $\}$ Quasi-optimal Prices
$P_g^* = £323$ $\}$ given if on £1500 Adv/man
$P_{1500}^* = £ \frac{156}{0}$ or if on £1500

Dependent upon the decisions of scheduling n/t o/t and
the number of value of C (production cost/item) the physicist
might to find P_0^* and derive the appropriate selling price for
Profit maximisation.

Fig 3.7 Embedded Normative Model (Mini Game)

PREDICTED RANK (BASED ON AIDS QUALITY)

ACTUAL RANK	S_1	S_2	S_3	S_4	S_5	S_6	S_7
1	1 2	1	1	1	1	1	1
2	2 1	2	3	2	1	2	3
3	– 1	1	2	4	1	3	1
4	3 (S) (B)	3	1	1	1	5	2
5		4	1		1	4	6
6		5	1	3	1	6	4
7		1	1		1	1	6
8		1			1		7
AIDED CoS	5/6	5/8	3/8	4/6	8/8	6/7	7/8

Fig 3.8 Consultant's Aid Success in Maxi Game

Cattell 16PF Scales

FACTOR	— Range	+
A	Reserved	Outgoing
B	Less Intelligent	More Intelligent
C	Emotional	Stable
E	Humble	Assertive
F	Sober	Happy-go-lucky
G	Expedient	Conscientious
H	Shy	Venturesome
I	Tough-minded	Tender Minded
L	Trusting	Suspicious
M	Practical	Imaginative
N	Forthright	Shrewd
O	Placid	Apprehensive
Q1	Conservative	Experimenting
Q2	Group-tied	Self-Sufficient
Q3	Casual	Controlled
Q4	Relaxed	Tense

[IQ/g | Dull | Cognitive, Bright..]

Fig 3.9 Simple Trait Descriptors (Cattell 16PF)

(68.) Group A: NOT AIDED
    I   Worldly, self-assured, lazy.
    II   Reserved, prudent, cautious
    III Staid, tough-minded, practical.
    → (Generally: self-assured, inward looking with good "cognitive skills")

(?)   Group B: PARTIALLY AIDED
    I   Self-reliant, lazy, ready to venture out
    II   Suspicious, less intelligent, fretful
    III Calculating, worldly but cynical.
    → (Generally: Worrying, suspicious but not quite bright enough to venture out of situations even though willing.)

(45.)   Group C: FULLY AIDED
    I   Stable, socially bold, leadership skills
    II   Low team integration, opinionated
    III Bright, high cognitive skills, assured
    → (Generally: Intelligent assured group knowing their individuating aspects & "creative" enough to integrate for the task)

Fig 3.10 Brief Results of a Factor Analysis
        quoting first 3 of 6 factors founds by oblique
        rotation.

"The Past and Future in Decision"

Broadbent : Well, this title is a little bit misleading. The talk is
really a continuation of the things I was saying last year. Those of
you who were here will remember that I was talking partly about some
lab experiments, but also partly about looking at various business
games and drawing anecdotal conclusions from looking at them. One
of them was the game that John Kidd talked about earlier and the
results from all of them were rather similar to those that he was
reporting this morning.That is, problems arising in situations which
were, at least in principle, completely determinate. Problems particularly
where there were lagged effects and overlaps in time between decisions
and the effects which made the latter hard to predict. And thirdly,
some conflict between acting in the situation and learning about it.
Staying alive while you find out what goes on.

What I want to say today is mostly an attempt to get a bit further in
this by looking at data from one particular business game, also carried
out at Henley at the Administrative Staff College. This game happened
to be particularly useful for our purposes because there was no
competition involved, so any team would give you results that weren't
affected by the decisions of the other teams, and this happened to be
very useful. There was also a certain sort of wry timeliness about it
because what the teams were doing in this case was controlling a computer
model of the U.K. economy. Every simulated three month period they had
to take three decisions. One of these was to set the level of the
Government expenditure which is known as G. One was to set the marginal
tax rate which was R; we didn't discriminate between different kinds of
taxes or anything like that. The third thing was to set a restriction,
if desired, on the amount of the money supply, which is known as K or
in fact, the restriction is known as K*. (Now these symbols are
conventional amongst economists. They are not completely arbitrary.
K is used for the money stock because M is already used for imports
and M is used for imports because I is already used for the amount
invested in the economy. You just have to learn these things I'm afraid

and stick to them).

Now the people were, as in the games I was talking about last year,
middle aged British managers in teams of 3 or 4. They had had some
lectures on economics and they were presented with the information
that they were going to have to take these decisions, and with a
print out of the performance of the economy for each of 8 three month
periods previously. In addition to the 3 control quantities which
were set for each of the three month periods, they got a print out of
large numbers of dependent variables, such as the level of personal
consumption expenditure, the personal disposable income, the retained
earnings of business, the value of the money stock, the redemption
yield on short term investments, the total taxes, etc etc. In amongst
all this stuff that came pouring out of the computer, they got certain
quantities that might well be of some interest, and importance, like
the proportion of the work force unemployed, the rate of inflation,
the value of the gross national product, the amount of exports, the
amount of imports, and it was suggested to them in the period that I'm
talking about that their object was to maximise the well being of the
society. The well being of the society is defined as twice the square
root of the gross national product, minus the square of ten times the
difference between the unemployment rate and two, minus the square of
half the inflation rate, plus one tenth of the balance of payments
surplus, if any.

You will be glad to hear that as a result of researches this
instruction is no longer given.

Anyhow, they were trying to maximise that; they looked at the previous
eight periods, and then they went in and they took these decisions at
the rate of about an hour and a half per decision. They went on
usually for five or six periods, the whole thing taking three or four
days. And this has been done for some years at Henley. (The model of
the economy they are working with I should explain, was constructed by
economists in the Staff College, particularly Harry Slater and the
Principal of the College. I have the equations here. There are only
22 equations. It's a very simple model. And there's no great problem,
I understand, in the computing side of it. In fact, it could be done

on a lab computer, but it was actually done on line to a time sharing
system using teletypes).

People experiencing this come out with a great feeling of having learned
something. You know, it is a great learning experience. And if you
say, well what have you learned, they say, well we now understand
economics or something, and we will buy Penguin books about economics
in future. So we felt perhaps it would be nice to find out what people
have learned as a result of this experience. And so we gave them a
questionnaire which was the simplest possible questionnaire of economic
knowledge. We said to them, suppose that you increase the level of
Government expenditure, what will the effect of this be on each of the
following three quantities, other things being equal?

$$
\begin{array}{cccc}
 & G & R & K^{*} \\
U & & & \\
 & \downarrow & \uparrow & \\
Y & \cdot & \cdot & \\
F & \rule{4cm}{0.4pt} & &
\end{array}
$$
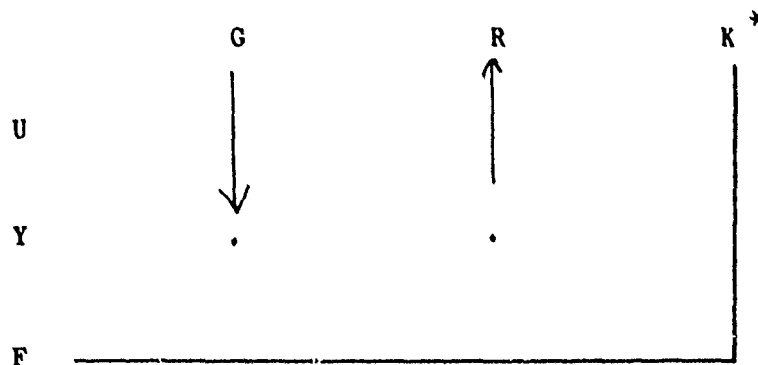
One of the dependent variables was the level of unemployment. And one
was the level of the gross national product. Now the gross national
product is Y. I don't know why, but it is. And the third question,
what happens to the rate of inflation (which, of course, is F because
I is imports, remember? Now, it isn't, is it? I is investment, that's
right. M is imports). And we asked them, suppose you raise the level
of Government expenditure; as a result, is unemployment going to go
up, down, or stay the same? Is the GNP going to go up, down, or stay
the same? And is the level of inflation going to go up, down, or stay
the same? We asked the "which" question about all the nine cells, every
combination of independent and dependent variables.

We handed out the questionnaires in pairs, each pair bearing the same
number; we gave it to a manager and said, fill one in and give it back
to us before you play the game and then fill the other in and give it

back to us after you've played the game. By this means nobody will know who filled in the questionnaire, but we will know which questionnaires are filled in by the same person. So these distinguished people, many of whom hold high ranks in British industry, would fill this in quite honestly, you, without feeling that people knew what they were giving away.

In fact, quite a large number of people dropped off and didn't give us back the second one. We ultimately got 30 managers to fill in before and after questionnaires of this type, where they had played the game for the best part of a week inbetween.

Now we discovered before they went in that most of them know that if you raise Government expenditure, you reduce unemployment. And most of them knew that if you raise the tax level, you increase unemployment, and similarly, the effects on the GNP. Nobody knew really what happens as a result of restricting the money supply, nor did anybody know anything about the causes of inflation. So that down in the cells linking those variables, there were very small numbers of people getting the right answer. In the better known relationships, 80 to 85% would get the right answer, but for the unfamiliar ones, somewhere in the sort of 5 to 20% region. They then experienced playing the game, and afterwards they came out and they filled in the same questionnaire, and they were really not much worse afterwards than they were before. They still knew the relationships they understood before, they were a little less sure but they weren't significantly impaired. But they were not improved on the relationships they didn't know; indeed, they were, if anything, a little worse. So apparently they did not, by playing this game, learn verbally what the effects of their decisions were.

We did have an extra question which was, are any of these effects only operative after a lag? And there was a just significant increase in the number of people who said, yes, the effects on inflation only happen after a lag. Even at the end of five periods, that percentage was less than 50% so on the whole they had not learned about that and they certainly hadn't learned about any of the other things. The ones who

learned this were quite right; if you took any decision whatever, it
had no effect on the inflation rate in the next three months' period,
and, furthermore, when it did have an effect, what it changed was the
rate of change of inflation. So anything sizeable that appeared in
inflation appeared only after this rate of change had had an effect
for several periods, so that when you took a decision you were actually
altering the rate of change of inflation after a constant lag. And
they had just significantly learned this. It was quite surprising that
very few of them had learned it, because everybody's results were
available to everybody else and in the first period, after all the
teams were launched, they all got identical inflation rates despite
the fact they'd all taken different decisions. So they really ought
to have caught onto that one, but they didn't.

These results then suggests that people have not learned intellectually
very much from this experience of interacting with the system. From
talking to both pupils and staff it was felt that this questionnaire
might not be the right kind of measure, because possibly what mattered
was the quantitative size of the relationships, not the direction, but
the quantitative size. We had two batches of managers, making fourteen
teams in all, from whom we had got with every decision a prediction as
to what was going to happen as a result of that decision. And we got
the predictions, therefore, about GNP, about inflation, about unemploy-
ment. Let's consider particularly GNP and inflation rate.

We treated the two courses separately statistically, but both gave
similar results. I'll just give you the results for the second course,
largely because they are funnier, but the significance of the results
is just the same in both cases. First of all, we take the prediction
of GNP before the very first session. The first question is, how are
we going to score the prediction? Now I'm sure lots of people would
have different views on this. What we actually did was take the
percentage error in the amount of change from the previous value. I
think anything one did would be a little bit arbitrary, but that is
corrected then for changes in money values for instance, which you are
going to get into trouble with otherwise, if there's an inflation. So
we did that. And we found that on the very first trial their

prediction as to what GNP was going to result from their decisions was in error by $147\frac{1}{2}\%$, so they weren't predicting terribly well.

	Y	F
First trial	$147\frac{1}{2}\%$	7%
Last trial	15%	1,152%

However, when they had played this game for the best part of a week, their prediction on the last trial was in error by only 15% and, therefore, this difference is significant. (In this course and also in the other course). There is no question that interacting with this system teaches you quantitatively what is going to happen to GNP as a result of your decisions. We also asked about inflation rate, and on the first decision this particular course was very good. They were only in error by 7%, and on the last decision they were in error by 1,152% but this difference is not statistically significant, it's due to one or two groups. However, that's why I gave you the second course first. The other course similarly gave no significant difference in inflation rate.

So the result of this week of effort is that people are no better able to predict inflation by the end of all these trials on the system than they were at the beginning. At this point I was going in and out of Henley under an armed guard! People were saying, we've gone to all this trouble to set up this game and teach people how to do this, and you're saying we're not teaching them anything. So I said, oh no, I believe it, you are teaching them something, but I don't quite know what it is, that's all. And we, therefore, resorted to the last refuge of scoundrels, which is individual differences. We started looking at the differences between the different syndicates. Now this was, in fact, quite interesting. Suppose that you intercorrelate the performance of the various syndicates on different measures. You can take, for instance, the GNP, the balance of payments surplus (exports minus imports), the inflation rate, and unemployment. The pattern of correlations was the same in both courses. Suppose you take for one particular course the performance at holding down unemployment and the performance at increasing the GNP and you do a rank order correlation across the various teams. You find there is a correlation of 1.0, the

best team at controlling unemployment is also the best team at keeping
up the GNP. And you say, well, you know, some teams are just brighter
than others aren't they? And indeed, it does look rather like that,
because if you take the best team at maintaining the balance of
payments and the best team at holding down inflation, the correlation
is not quite perfect, but it's 0.9. They are very closely associated.
But, when you look at the other correlations, you find that they are
all negative.

	Y	X – M	F
U	1.0	–.52	–.43
Y		–.24	–.43
X – M			0.9

In other words, teams are good at controlling unemployment, or they
are good at controlling inflation, but they are not generally good or
bad at controlling both.

Zeidner : What is there a significant correlation?

Broadbent : Well, there are seven teams, so the correlation of
unemployment + GNP is significant and the correlation of balance of
payments + inflation is significant, but I am not so sure about the
negative ones. The positive correlations are significant. The negative
ones won't be significant on a sample of 7. But the reason I keep the two
courses separate, is that during one course, although there are 7
teams, they had exactly the same lectures, whereas on a different
course, the lecturer may have changed his technique slightly. So I
didn't like to correlate across courses and make n = 14. But the
replication of the negative correlations on two samples each of 7
does at least make us confident that the relationship is zero or
negative. This begins to suggest that we've got some teams that
are good at unemployment and some that are good at inflation. You
could then say, "well maybe that's change; they go in and they make
a random decision and the characteristics of the economy or the model
of the economy are such that a random decision will either give you
good unemployment and bad inflation, or vice versa". So you begin
to wonder whether they are in fact trying to achieve good
unemployment or good inflation.

The first thing you do is look at the predictions they made on the
very first run. Suppose you take each of the four dependent
quantities, rank order them, and then add the ranks on GNP and
unemployment, add the ranks on balance of payments and inflation
and subtract the total of the latter from the total of the former.
You've got a rough measure of whether this is a team that goes for
unemployment or this is a team that goes for inflation. So you
calculate the mean for performance on the last run and then you do
the same thing with their predictions on the first run, before any-
thing has happened at all. Now their predictions are related to their
performance. You find in the first course that the correlations is
.5 and in the second course that the correlations is .43, so they are
on the whole predicting the bias that they get in their results.

Well, at this point one begins to say, we would like to measure the
utilities of an individual team. That's where we ran into trouble.
We thought it would be nice to do this because of one of the most
interesting things in political thinking that's happened here lately.
Those of you who keep the journal of the Royal Statistical Society
as bedside reading, will have noticed a paper by a socialist member
of parliament called "Control of a Noisy Economy" and this is by
Jeremy Bray, MP. What he is saying is that when people take decisions
about the control of an economy, they tend to get a forecast as to what
is going to happen, and then they take a decision which is based on
the forecast, and off they go. Now the one thing that is certain
about this forecast is that it is not completely accurate because
there is noise in the system. Therefore, they know that all the
controlled quantities are, in fact, going to deviate from the
trajectory that is predicted. Once they have deviated beyond some
level that is politically unacceptable a fresh decision will be taken.
As soon as that fresh decision is taken, the assumptions of the
original prediction cease to be true. Therefore, says Bray, the
decision at any point should not be in the form of "if we take a
decision now and sit back and do nothing, what will happen", but
rather, "if we set control parameters to certain values, what
trajectory can we maintain within certain limits", and he shows that
if we had done this ever since 1945 we would, in this country, have
done better. This is a very interesting idea. What he is suggesting

is that the Minister, when a government takes office, should be asked
to set out his utilities in a sophisticated and detailed form, and
then lower order control decisions should be taken into the economy
on the basis of these priorities. He supposes that it is quite easy
to obtain the utilities and preferences of politicians, since, after
all, they have all formulated political programmes in the past. I am
afraid we don't find it quite so easy. For a start, the teams are
all asked before they start running the economy to produce a statement
of their policy. Their policy is usually to hold down inflation so
far as is consistent with not having an excessive increase in
unemployment and meanwhile maintaining a balance of payments surplus
and increasing GNP. That would do for any of them you see. They
are very reluctant to state it in numbers. Now we only began to get
wise to this after the first course, and we haven't, therefore, got
data from them. On the second course studied, we started out by going
in with the welfare function that I mentioned earlier, and giving a
number of paired comparisons based on this function. You remember,
twice the square root of the GNP, minus the square of ten times the
difference between the unemployment rate and two, and so on. We did
paired comparisons on the question, which would you mind most, a rise
in unemployment from two to two point five, or an increase in the
rate of inflation from such and such to such and such?"

These questions were all based on the welfare function. He assumed
that some teams would have a marginal preference for allowing some
unemployment and some teams would have a marginal preference for
allowing certain amount of inflation. In fact, the welfare function
was quite clearly completely wrong because every team always went over
to one side. In fact, they all disliked inflation much more than this
function says. So the function was clearly nonsense. Also we weren't
getting any variability between the teams so we couldn't do any
correlations for their performance.

We did just get one last question which seemed to show something.
We put in one open ended question which didn't rely on the approximate
truth of the welfare function. We said "Which of the things we've
given above (which were all unpleasant and involved rises in inflation

and rises in unemployment) would be worthwhile if it was accompanied by a 10% GNP increase?" Five of seven syndicates filled that one in. The other two said no, we're not prepared to answer that question. We don't think we want inflation to go up regardless and we just have to find some other way.

Zeidner : Which kind of team were they?

Broadbent : Well, it's interesting you should say that because when we finally took the bias score that we were using here and we looked to see who these two teams were, they were the ones who wanted to control inflation and let unemployment rip. They weren't prepared to say so. Remember that there's a very large correlation between controlling unemployment and keeping up the GNP.   If you keep up the GNP you keep up the profits of your company.   My experience is that these managers who have the reputation of being right wing, when you get some rough measure of their utilities actually care more about avoiding unemployment than, say, apparently left wing academics who don't really have this sort of intimate concern with the problem. Anyway, we only had five syndicates left, so I'm not going to claim that this correlation is significant, but we could relate performance on the first trial to our measure of utilities. That is, the bias on performance and the bias utility estimate. They correlated absolutely zero on the first trial, but by the last trial the correlation was 0.7. Now the main point of these results, apart from the problem of getting utilities, is that we've got another problem. The decision makers are able, apparently, to get performance out of this computer model which corresponds with what they want in some sense, at least to what they predict is going to happen. And yet they cannot give an adequate verbal account as to how the system is working. So how do they do it? I'm not going to answer that one, merely leave it as an exercise to the listener. But I think it fits in with a lot of the things that John was saying this morning, that somehow the heuristics that people use in interacting with the system like this don't correspond to anything you would see as a conscious systems analysis. Yet somehow they achieve something that corresponds to what they want to get.   I have one hunch.

I don't quite know how to test it, though I'm trying to think up things. But it does seem to me possible that when you have the actual experience of playing with a system like this, you build in to your nervous system each of the relationships that exist between the various quantities. Not 22 equations, I don't think, when you only have five trials or something, but at least some of the relationships. Then when you are taking a decision, you do it on the basis of a parallel processing or analogue type computation for these several relationships which then cannot be given as a serial verbal description, so that you're turning yourself into an analogue computer and are unable to get a serial print out. And if that is so it raises, of course, terrifying possibilities for everyday life. We all know the kind of manager who says, "well, it's all very well you young chaps with your book learning, but it's my experience of the situation from years past that your ideas won't work". He may indeed be right, provided that the present situation still obeys the rules, that his experience suggested. Of course, all too frequently it doesn't, but then in that case, maybe you have to train people, say, with a computer model, until they have built in all the factors in the situation and then you can go out and try them on the new situation.

Well, that is really all that I have on that piece of work, but I did want also to mention one other thing fairly quickly about another aspect of spreading in time. Can I go back and link this up with the title? The reason these results concern spreading in time is that this inflation rate thing is a lagged effect. I should say again that particular relationship is the one of key importance.

$$F_t = F_{t-1} \left[\left\{ \frac{Y_{t-1}}{Q_{t-1}} + 0.06 \right\}^{\frac{1}{3}} \cdot \left\{ 1 + 4 \; \frac{K^*_{t-1} - K^*_{t-2} - 0.04)}{K^*_{t-2}} \right\}^{\frac{1}{3}} \cdot \left\{ \frac{F_{t-1}}{F_{t-2}} \right\} \right]$$

There is a rationale for this model, although I think one could establish this sort of equation purely empirically. If the rationale is valid, it's based on the argument that the setting of prices and of wage

bargains and so on, changes at a rate that is determined by the rate of change of prices in the previous period. People start incorporating into their calculations things about the rate of change in the previous period rather than the absolute values. And that's a sort of psycho-economic assumption. One hardly dares, in liberal circles anyway, to mention where this idea comes from because it's associated very often with suggested remedies which I, for one, would not support at all, but it is, of course, Milton Friedman; and my impression, talking to economists, is that most of them do feel that this relationship is right, given our current system. This is what happens and that, of course, will be why inflation rate tends not to be noticed at first and then suddenly turn round and kick people because it comes in as a lagged effect in a control problem of this kind. So the problem of economics is a matter of learning to control the system when the effects are coming back at you well after the decisions; and this is very difficult as John was saying. If we are going to work with a system, where you've got to remember the past in order to do your present decisions, then think ahead to the future on the basis of extrapolating past relationships, then how you organise your memory becomes very important:

You remember I talked last year about different possible ways of organising a retrieval system and the way in which people possibly do so.

I was contrasting a system of hierarchical classification where you may have a memory for certain cities and these perhaps are divided into the American and the European cities, and then inside America you have US and Canada and, inside Europe, you may have UK and the mainland and, inside the UK, you might conceivably separate Scotland and England. I'm taking it more or less that this is the psychological map of the world.
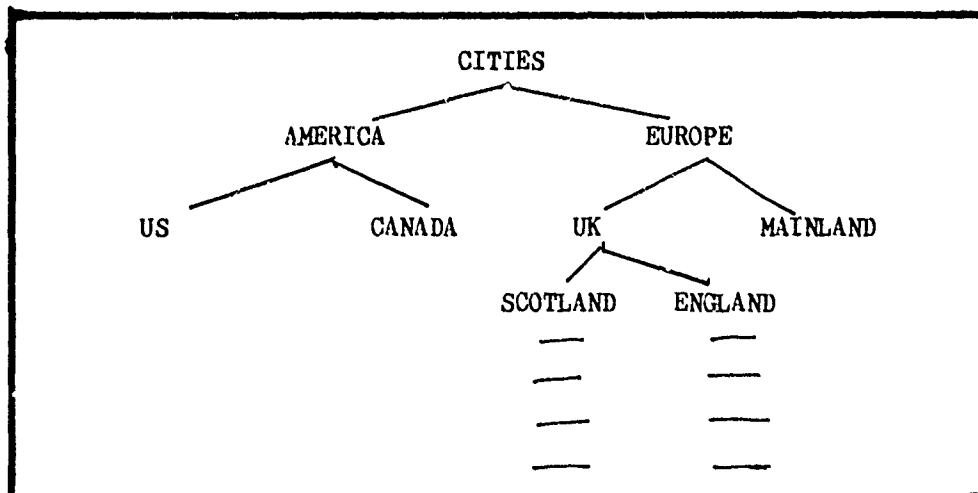
Broadbent : Anyway, you can draw out a hierarchical tree like this
and then at the end of it you can put in a list of items and you
have four Scottish cities and so on under each of the other categories.
Or alternatively, you can have a matrix type classification where you
have, say, cities or rivers that are American or European, that are
large or small, that begin with letters in the first part of the
alphabet or the second half of the alphabet and so on. It's difficult
to draw that because, of course, it will become a sort of Necker cube
with different dimensions dividing up different cells in the cube; and
once you get beyond three dimensions, it gets even harder. The point
is that it is possible to organise a retrieval system in either way
and I did give some evidence last year that different people do so to
different extents. Now we have been doing experiments in which you
impose these two structures on people and ask them to remember material
and I said last year that we had rather little difference between these
two systems. Both of them give apparently fairly equal average amounts
of memory. Now I want really to mention only one result this year which
we've now got by further analysis.

Well, we did a number of experiments trying to compare hierarchical
versus matrix type structures to see what happened and all of them on
average showed the matrix very similar to the hierarchy. That includes,
for instance, the situation we tried, for example, the same number of
items to be recalled with a very large hierarchy, having many branches
but only a few words on each branch; or a small hierarchy with few
branches but many items of the same type. And our prediction was in that
case that a large hierarchy would do better than a large matrix because
of the greater ease of remembering, say, a bottom branch as being
Scotland, since if it was Scotland, it also had to be UK and also had
to be European and so we thought that that would give a particularly
good arrangement for the hierarchical system rather than the matrix.
But there was nothing in that. Conversely we had cases with a very
small number of cells with a very large number of words in each cell
and we thought this would be particularly good for the matrix system,
since in each case you get two or more cues going into each cell and
that would be perhaps beneficial. Whatever changes of this kind we

tried, we could never get any differences between hierarchy and matrix.

And then we began to look at a new particular measure. We looked at a hierarchical tree and we considered whether, if the subject forgot one particular cue word like Scottish, what happened to his probability of recall for items, and the answer was that if he remembered the cue words (as it might be, Scottish), then his chances of remembering Edinburgh and Inverness are very high like .75, whereas if he forgot

	Hierarchy	Matrix
Cue Recalled	0.75	0.5
Cue Not Recalled	0.25	—

the word Scottish, his chances of remembering the other two words were very low, like .25, so that there's an enormous difference in the probability of recall, depending on whether you recalled the cue or not. Now in the case of the matrix of corresponding size the results are quite different. Firstly, you'd never get cases in which there is no cue word recalled. The reason is that every word in a matrix of 16 cells has got 4 different cue words all coming in from different directions and as long as you remember even a quarter of the cue words you will remember at least one that points to the thing you want. Therefore, you just never get cases in which there is no cue word. Secondly, your probability of recall given the cue word is not as high as the probability of recall given the cue word in the matrix case; it is only about .5. The effect is, therefore, that you've got the same average level of performance in the two conditions. What happens in the hierarchy, however, is that you remember some things perfectly and others not at all. Whereas in the matrix you remember everything rather sloppily. Now if we can extend this to the individual difference case that I mentioned earlier, it seems to me that we get two quite different patterns of memory for the past.

Hierarchy   Table 6

MATRIX
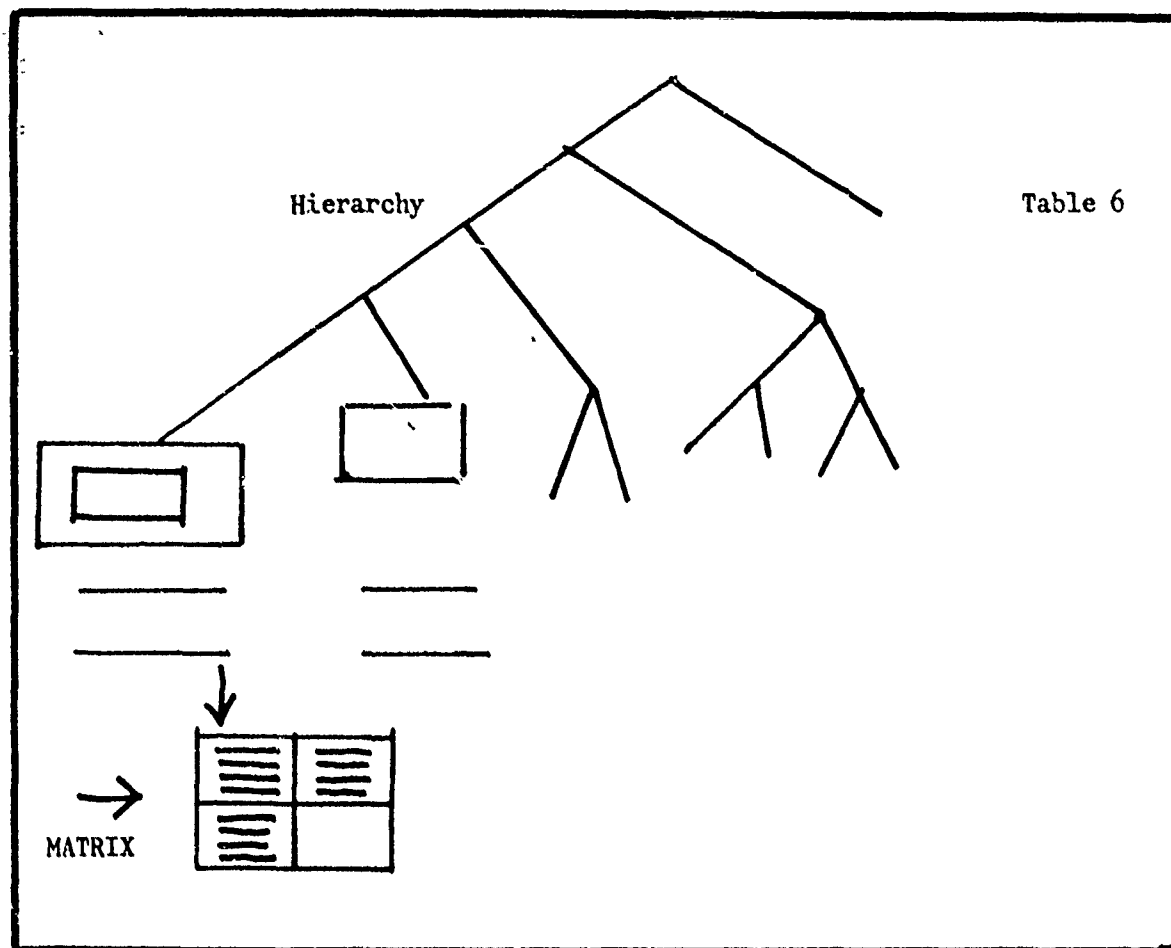
One type (hierarchy), you might say, is highly organised, bureau-
cratic and blinkered, a person who remembers some
things absolutely marvellously but is completely blind to everything
outside of the things that he's concerned with. While the other has
this complex rather robust network of associations in which he remembers
things with an occasional probability of error. The average performance
works out about the same, but the detailed structure of the performance
will work out very differently. Now I am afraid this talk obviously
breaks into two halves because I have no way in which I can say, well,
these are the chaps you see who don't care if there's unemployment.
And indeed, I don't think I would really predict that, but I do suspect
that when we can get at the individuals, more precisely, those who tend
to one kind of structuring will differ in control tasks as well.

This tendency to remember some things and then forget others completely
may well give a rational but verbal type of control, whereas the others
will crash through the middle, probably with the sort of analogue/
intuitive form of performance that I've been mentioning earlier; and
which of these would be better, of course, will depend on what it is
you are asking them to do. But those are, I am afraid, all the things
that I've got, and that's it.

NOTE : There was a discussion by Dr Pask and Dr Zeidner of the
questions posed in the Introduction.

Broadbent : My vote would be for saying there are a collection of skills rather
than a single one. Admittedly at a higher level you have to have an
integration of the separate functions, but you need the elements first.
A cricket player has to have a collection of skills rather than simply
being a good cricket player and nevertheless being a good player
requires a balance of the component skills. Now, the real bearing of
the things I've been saying on these questions is that if you have
people who have prolonged experience of certain situations, they will
build in some of these skills but such skills are not readily imparted
by verbal instruction. I think, as I was saying last year, that the
different retrieval strategies are probably of this kind and are
associated with type of education, science versus arts, things of that
sort. There may be some temperamental base but, purely on anecdotal basis,
I feel that the people who operate on trees are scientists on the whole.
And I suspect that this is just the result of years of going through
operations of this kind. On a shorter time scale, putting people into
interaction with the model of the economy seems to produce changes in
their behaviour. Whether it would transfer to other situations or
whether it just holds true in that situation is another problem, but
at least they seem to have incorporated or introjected something that
they can't describe. So on the whole I'm hopeful about training the
collection of skills that are used for decision making, in a particular
context.

Kaplan : You said that in long years of experience people build in skills
not verbally desirable. Would you care to speculate on what you think
each still did or how it was learned?

Broadbent : It won't be speculation if it can't be spoken, will it?
I referred to a collection of specific sub-skills. The sort of sub-
skills I would think of are at the very least something like an intake
from the outer world which represents the sensory environment, then you've
got some kind of short term memory, probably subdivided. (You know, I
usually say that I try and confine myself to five or six kinds of short
term memory these days). You've also got a processor and you've got a
long term memory. Then its almost analytic that each of these can be
fed by and can feed the others, and that you have to learn certain
routines (which are presumably in long term memory) for swopping
information from the sensory intake into short term memory, or possibly
into long term memory, or from long term into short term memory, each
as a response to a particular situation. For instance, you may become
conscious that you need, say, some toothpaste. You then have to put
into short term memory, I need toothpaste, I am going to get some.

Table 7
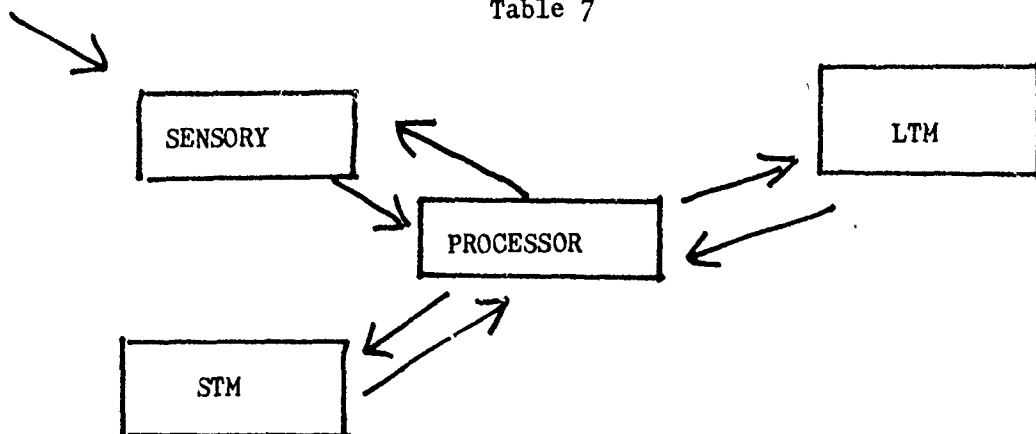
LTM = Long Term Memory.   STM = Short Term Memory

You go outside and you see the street. To take the decision to turn
right you then have to have in short term memory the fact that you're
going to a place where you can buy toothpaste, otherwise the stimulus
of the street might take you to a coffee shop. So the goal has got to
be in short term memory while you are executing that task.

You will probably have to go down a sort of branching tree. I mean, if
you are at this hotel you may say, well, there's some shops down the
hill, so, given a goal of toothpaste you add a sub-goal of going down
the hill, and then a sub-sub-goal of turning right to go down the hill.
As you don't, in fact, know that there's a chemist down there, the sub-
goal is heuristic, and if you like to represent the process instead of
the structure of the system, you would get a tree of sub-goals being set
up one below the other. Each of them is being put in, in turn, into the
short term memory and when its completed its ejected and another one's
put in. Now I've gone off into this digression because it's obvious
that a system which operates like this or an artificial intelligence
simulation of something operating like this, must have certain very
long lasting rules that are in long term memory. These rules, in a
certain situation, do something which may be to put a sub-goal in STM,
or maybe to act on the outside world or may even be to change things
in long term memory. But these long lasting routines that are in LTM
have to be acquired in some way and I don't think they can be acquired
just by telling someone about them. If you just tell someone about
something it goes in, it will interact with what's in long term memory
and it'll go into short term memory. It may perhaps remain in LTM as
a statement. But it will not alter rules of behaviour. You can say
to someone, Mr Callaghan has become Prime Minister, and it is, in one
sense, in long term memory. But he doesn't feel Mr Callaghan is
Prime Minister, you know, until he's operated on it a bit and thought
about it. When he has to act on some practical problem he will start
his old habits, and then say, "Oh yes, it's Callaghan now". Thus he
keeps operating on this until he changes the routines in LTM. The
sort of thing I'm talking about though is much more long lasting than
that. General rules of action are unlikely to be changed, and consist
of habits or heuristics which people adopt when they go into a new
decision making situation. John will probably agree that if you take
experienced managers and watch them playing a business game, you can
see the heuristics coming out. You plonk a brief in front of them
and a group of them sit around and they are all reading the briefs and
silence ensues for about five minutes. Then you can practically predict,
either as the first or second remark, someone will say, "what are the
things we have to decide? What alternatives have we?" And someone
else will then look and scan down the list and they'll produce the first

one and say, in the economics game, we have to decide what level of government expenditure to set. The next statement will be, "what happened in previous quarters, what is the past experience on this particular thing?" These are absolutely regular remarks at Henley.

Kidd : I'd like to suggest though that they have some perhaps some inbuilt algorithm for reduction of very simple binary searches. That they would be immensely overloaded if they had to think in detailed terms about the task.

Broadbent : I think that's right and I think it is a rational behaviour.

Kidd : But perhaps the great problem is that algorithm is not quite right for what they've now got to do and they've got themselves into an anchored process perhaps by wrong transparency.

Broadbent : Well, the point I was trying to make was just that if you take people who have survived 20 years in industry, they have these routines in LTM which they bring out in a new situation. In the class of problems that I've been faced with, and for most Henley people, the first thing is to say, "well, what options have we got?" And then it begins to make sense to them and then they say, "well, what did we do last time?" Maybe they won't do what they did last time, but at least it structures the problem and so what these people have learned are these strategies which are actually heuristic, I mean they won't guarantee them the right answer, but they make success more probable.

Kaplan : Do you think that if it were possible, in fact, to train in these heuristics as you call them which you just said you don't think you could do, would those constitute the sorts of things that people here have been talking about when they said, when they've spoken about learning to learn, or having procedures which could transfer, generalise, could be applied in a variety of situations.

Broadbent : Yes, I think they would and I wouldn't like to say that
I didn't feel they could be trained. On the contrary, I feel they
can be trained but they can't be communicated simply by one verbal
proposition. That's not it. You've got to use them, live with them,
interact with them. I'm getting, of course, way off, into speculative
intuitions about what's going on here. It certainly isn't implicit in
the data we got here.

Brehmer : I was wondering about the empirical basis for your statement
that these managers couldn't really tell what they were learning. You
have this table with the G, R, K and the other letters. Weren't there
any individual differences, and were these individual differences
correlated with some other behaviour?

Broadbent : Well, unfortunately because of the anonymity we couldn't
correlate the forms. Also it was team performance which is another
problem. But we wondered very much about this. However, for effects
of tax on unemployment, well over 80%, I think 85 to 90% of the people
who returned questionnaires knew that. Whereas for effects of money
supply on inflation it's under 20%, probably only about 10%,; so the
range of individual differences can only be very small.

Brehmer : Yes, but the correlations among teams seem to be quite
small too. Certainly, there are possibilities for some teams behaving
in a quite different way from others. And that's a fact that might be
something additional to what they are reporting verbally. Whereas it
seems somewhat unlikely to me that they wouldn't get the direction right.
They might not get the slope right.

Broadbent : They are getting the direction right in the effect of tax
and expenditure on GNP and unemployment. I think what happens is that
most teams do not make proper use of restriction of the money supply
which, on this Milton Friedman model, will give them very good control
of inflation. They make use primarily of those two variables of which
at least a minority of them do know the direction of the effect. I
suspect myself that what they could verbalise about this situation is
that they know they can produce an effect on unemployment by changing
these variables, and they have a vague sense that unemployment is

somehow negatively related to inflation and by shooting for a higher
level of unemployment they will get a lower level of inflation. And,
therefore, what they do is transform the problem into one of trying to
achieve a target value of unemployment and that target may be higher or
lower, depending on their relative values, but they can't predict what's
happening with inflation.

Brehmer : That also means that they have problems in combining the
settings on the different variables, which is something you do not
test with that kind of device.

Broadbent : No, that's right. If they do go for unemployment I have
one piece of evidence which actually confirms this strategy specifically.
We had another group who went through and for whom I haven't reported
the data because there were only 3 teams, so it's simply observational
and anecdotal reporting. They had two trials and then they had another
lecture which no previous group had had. In this lecture there was
presented the Milton Friedman revision of the Phillips curve, which
relates not inflation but rate of change of inflation, to the level of
unemployment. A plot was shown for the 8 previous periods of the actual
obtained inflation values for different levels of unemployment; this
showed that the curve crosses the zero line somewhere between $3\frac{1}{2}$ and
4% unemployment. Which means that with this model of the economy, if
you can secure 4.1% unemployment, the rate of change of inflation will
go negative and you then sit tight and wait for everything to go right.
If you watch British Government policy, that is exactly what they are
doing. And it's working and the rate of change of inflation has been
coming down and that's why they've been feeling happy. The inflation
keeps going up but the rate of change of inflation is going down and
it's now gone over the top and is beginning to come down again, and if
they just hold the unemployment slightly above this key value, it will
go right down to the bottom.

Zeidner : Donald, what was the reaction,   of your subjects, to this
experiment? Were they pleased with it, did they find substance to it?
Were they co-operative? Their attitude.

Broadbent : Well, they varied enormously, of course. This is the
trouble. They mostly treated it as an interesting intellectual game,
I think one would say.

Zeidner : And they got involved? They took it seriously?

Broadbent : They got involved, oh yes. And indeed I haven't mentioned the point, of course, that during the period I mentioned they were still using this fantastic welfare function, as it was called. This was being printed out so they knew which team was doing best. Of course, what many teams were doing was trying to win the election, that is, they were trying to get the highest value of welfare in the last period of the game. With disaster happening the period after that they were certainly competing and attempting to do well.

Phillips : This theme that keeps coming up about not being able to find out what's going on by asking people the notion of fuzzy heuristics not capable of being verbalised. This is a puzzling one to me because I work in an institute where there are a lot of sociologists and social psychologists who use a process that they call social analysis to learn more about organisations and this is a very talky interactive kind of process where they go and talk to the people, and as far as I can make out, it's partly a kind of consciousness raising exercise. That in a way you're going in and giving them the tools with which to view their own processes within the organisation, when they can then feed back to you on a subsequent occasion what they have learned and then you feed back to them your impressions of their impressions, and over a period of time both you and they seem to learn more and more about the organisation and particularly the differences between the latent and the manifest structures of the organisation. I wonder if our emphasis on being good experimental psychologists where we kind of do a before and after measure and don't muck about at all inbetween, doesn't lead us to find that this isn't a very useful approach?

Broadbent : Well, I did ask a question at the briefing, not on this, but at a debriefing session where we reported that indeed they had been giving performances comparable to their predictions, and I then said, was this consciously formalised? To which someone immediately said, yes, we knew that if we raised the level of unemployment we were shooting at we would get a lower inflation rate. I avoided prodding "what do you know

about the form of the equation", and I think there's a level at which it
is conscious and that it involves verbal behaviour, and yet it isn't
conscious in that they can't put numbers on it or say that it's a linear
relationship, or whatever the shape of the function is.

Pask : A point that might be brought up in this connection is that
there's a difference between describing a relation, a particular dynamic
relation of which this is one, with lags in it and so on, and being able
to compute it, which these rather cack handed equations presumably would
allow us to do. I couldn't remember a lot of them but that would only
be one way of computing it.

The fact is that sometimes recall in the sense of being able to describe
is adequate, sometimes not. It depends on whether I have to mentally
compute a relation or not and if so, upon whether I am able to build
up a mental computation from a description of what is to be done.

Almost unlearned skills are one case where I have the co-ordinants of,
say, GNP and can say GNP is related to some other quantities, whereas
I can't say how its related. I couldn't give a procedure for calculating
or manipulating GNP, or a mechanism underlying that or an algorithm.
In fact, if I am a manager I am asked to do this (because managers do
other things than pontificate) and would have, at last, to construct an
algorithmic mechanism of some sort, which is learning. Having learned
I could say how to manipulate GNP.

The other extreme case is an overlearned skill. Here I can do the job
but cannot say how I do it because there are so many algorithms (mechanisms)
that the computation is fuzzy and I can't externalise it because I didn't
have one or a few ways of doing it. At best I could reconstruct (recall)
a representative procedure (I think probably Larry's points about fuzzy
computation are relevant here).

Most managerial decisions seem to be between these extremes. Presumably
management gaming could be used as a non verbal conversation for
externalising procedures, algorithms or whatever; in this case the games
are able to distinguish between ability to describe and not do, to describe
and do (but not say how) and the intermediary cases.

Another comment refers to transfer and style, or memory style (in this case, recall style). Stylistic differences seem to be curiously permanent, I don't know why. For example, people that have educational backgrounds seem to generalise style over everything academic and I imagine managers generalise over everything managerial, maybe they play tennis differently or go to the bar differently, but over the particular area of competence the style seems to generalise in a quite unaccountable way.

Finally, somebody brought up "learning to learn" and I would merely point out that probably transfer and differences of style are not incompatible. There are people who can probably use a bit of many styles of learning, recalling, doing and maybe they can be trained to do so, but "learning to learn" means rather more than that. It means (I am asking a question here, is the following true of management as it would be of education?), not only knowing what style you adopt best but also an ability to give a structure to an environment which would normally come to us unstructured.

SECTION 4

4. Application of Specific Decision Theoretic Methods

Overview : Two papers, one on "Portfolio Theory" presented by
Dr Borcherding and one on "Praxeology" by Dr Schaefer are included
in this section.

So far as decision is concerned, Dr Borcherding came up, after meticulous
experimentation, with a negative result. People cannot predict random
events, nor is it likely that the input assessment data required by the
portfolio model could be obtained reliably in systems of the type under
consideration. Since Dr Borcherding was unable to self-edit her paper
in the interval allowed, the abstract only is included, augmented by
summary comments. Her results are filed with the full transcript.

Dr Schaefer's paper is an account of the development of a science,
"Praxeology" which is concerned with rational or efficient action.
"Praxeology" very closely resembles Cybernetics and General System
Theory and it seems to have an isomorphic history though, for various
reasons, the ideation is independent (Praxeology is little known outside
Poland and most of the literature is in Polish). In his paper,
Dr Schaefer notes that Praxeology, until recently, had no real
foundation of its own and designated a collection of methods and
descriptions held together by faith on the part of its exponents:
partly engendered by pragmatic success. Gasparki and Szaniawski (two
authorities on the subject) agree to the lack of a formal basis but
Schaefer quotes them in respect of Decision Theory.

> " It has frequently been pointed out that the classical
> paradigm of decision theory - strategies, states of the
> world, utilities - has been exploited to such a degree
> that it has almost outlived its usefulness. Some radical
> change of the paradigm is generally felt to be necessary. "

Schaefer remarks on this quotation, in his abstract :
"(a) Strategies, states of world and utility are some of the primitives,
or basic notions of statistical decision theory. But: Is the theory
called 'Decision Theory' a paradigm - or is a situation such as 'bookbag-

and-pokerchip' a paradigm. This special experimental situation, maybe others as well, has been under criticism for a while. It is felt that Decision Theory should address itself to more real-life, more complex problems. Nevertheless, if it is possible to structure the situation in such a way that the Decision Theory approach is applicable, why not use it? Its methods are highly formalised and more or less easily applicable.

(b) The more serious question is, in my eyes, whether we have developed sufficient models for other situations, especially so-called ill-defined ones. This may be questioned. But some possibilities seem to arise: the system dynamics approach and fuzzy set theory, to name just two that come to my mind.

(c) Gasparski's and Szaniawski's criticisms lose some of their sharpness, but also of their relevance, if all recently developed models of Decision Theory are adopted for Praxeological use. Multi Attribute Utility Theory, bootstrapping, are but few examples."

"An attempt to use portfolio theory as decision aid"

Borcherding : Portfolio management is a complicated process
characterised by three steps: 1. security analysis, which means
inspection of individual shares, or, more generally, prospects,
which is an art, 2. portfolio analysis, ie. inspection of a
collection of shares, and making a theoretically based analysis of
portfolio composition and finally, 3. portfolio selection, the
decision to compose a specific portfolio.

Roughly sketched, the Markowitz mean-variance portfolio model
proceeds as follows : Exploiting only some specific characteristics
of the securities, a set of efficient or non-dominated portfolios are
determined by either solving a set of equations or by mathematical
programming techniques. Knowing an investor's utility of wealth,
a specific portfolio can be identified as the 'best' portfolio for
the decision-maker in question. It will be shown that the mean-
variance approach to portfolio selection supposes a quadratic utility
function.

To be able to actually use portfolio theory as a decision aid, the
decision-maker must specify some specific inputs. My main interest
was to empirically check whether people are able to deliver these
inputs. The main input is the probability distribution over future
security prices. For ten different securities, various groups of $\underline{S}$s
(experts in the normative/substantive sense) had to give subjective
probability distributions for future share prices: 1,2.., 6 months
into the future. This procedure was repeated six times, one month
apart. Subjects' assessments were analysed according to different kinds
of goodness.

Four groups of subjects were employed. Laymen, who are psychology
students; economics students; faculty members of the economics
department, who are specialists in and teach banking and assets;
statisticians who serve as substantive experts.

"Can Praxeology Help Decision Analysis?"

Schaefer : Together with many colleagues at the Darmstadt conference on "Subjective Probability; Utility, and Decision Making" (September 1975), I heard a lecture by Dr Gasparski on "Praxeology in relation to the decision analysis". The subject is not well known and some effort was needed to determine the background.

I shall first give a preliminary definition of what Praxeology is, or what it might be: second, a very short history of the word and the concept: thirdly, a more general characterisation of Praxeology, with special emphasis on its main fields of application and the techniques used by Praxeology or Praxeologists: fourthly, an account of the key concept, namely, efficiency; fifthly, a consideration of Praxeology in relation to some other branches of science and finally, I shall come back to "Can Praxeology Help Decision Analysis?"

First, Praxeology is defined as the science of rational action or as the science of efficient action. These are the characterisations given by all writers of Praxeology, including Kotarbinski, Oskar Lange and others.

The concept of efficiency plays a central role. Praxeology as a scientific effort is concerned with organised and goal directed human activity and the most relevant measure of goodness for this kind of activity is efficiency or rationality. The domain of Praxeology is not the classical scientific domain, but is concerned with all kinds of acts, for example, in engineering, management, economics, psychology, and so on, since it tries to define directives, develops measures for evaluation and prescribes how work should be done or even more often, how it should not be done.

In principle, Praxeology is indifferent to the ultimate goals of actions, may they be constructive or destructive or what characterisation you might ever find. Praxeology is a science of means, not of ends. The question is how can we attain a once defined goal, whatever the reasons for its definition might be. This is what I call

the 'technological' aspect of Praxeology. Now for writers who use the
word and the concept. According to a Polish author (Pszczolowski, 1966)
the word Praxeology was first used by Louis Bourdeau in 1882 as a name
for the science of functions or activities. Bourdeau distinguished
analytical Praxeology from synthetical Praxeology. Analytical
Praxeology is subdivided in two domains; elementary analytical
Praxeology or somatology (the science of functioning bodies), and
general analytical Praxeology or psychology. Synthetical Praxeology
is the science of relations between functions; again two distinctions
are made, comparative Praxeology and general Praxeology. As seen by
contemporary Polish authors, Praxeology is contained in comparative
Praxeology. Oskar Lange, the Polish economist and econometrician,
credits Espinosa an early French socialogist, with introducing the
term. Espinosa referred to Praxeology in 1890 and 1897 as the science
of most general kinds and principles of action in a world of beings
which are able to move. This restriction seems to be important in
this context. Another very early writer is Charles Arthur Mercier in
his treatise, "Conduct and its disorders".

The most prominent proponent of the emergent science of Praxeology was
Tadeusz Kotarbinski who wrote his first important paper in 1913, and
several later ones. His most fundamental work appeared in 1955 and
was translated into English under the title "Praxeology, an introduction
to the science of efficient action", and the English translation in 1965.
Independently the Polish Russian mathematician Slutski (sometimes
written Slucki) introduced praxeological considerations into the
realms of economics in a German manuscript, "Ein Beitrag zur formal-
praxeologischen Grundlegung der Okonomik" (the Ukranian Academy of
Science in 1926). As I learned from Gasparski, a complete history
of Praxeology is given in a book by Jean J. Ostrowski.

Finally, the Austrian economist Ludwig von Misses(brother of a famous
probabilistic thinker, R. von Misses), has written a very important
book on political economy (1940 in German, the definitive English
edition in 1949). Von Misses conceived Praxeology as a logical
science in line with logics or mathematics, emerging from political

economy. Von Misses' views are especially severaly attacked by Polish writers, for example, Oskar Lange. More recent contributions to Praxeology come from writers of different disciplines; for example, management science, systems theory, and decision theory. I will come back to these considerations later.

So far as I can see Praxeology is mainly developed and used in Poland. There is an institute for Praxeology similar to a German Sonderforschungs-sbereich (special research institute or unit), but outside Poland Praxeology and the concepts of these authors are not very well known. Inside Poland they are quite famous. They are well recognised, they are inter-disciplinary and as I will come to later they play quite an important role in the national economics because they say they can develop methods for giving guidelines to efficient actions. Since, in a socialist country, there is a central plan, they play an important role in constructing this national economic plan.

My third task is to give a slightly more general characterisation of Praxeology. The intention is to formulate general assertions on the basis of efficiency, which apply to all kinds of purposeful and goal-directed actions. One important goal of Praxeology is to formulate directives or general instructions for efficient action and also to produce hints and warnings to avoid common failures and errors. To achieve these goals it is necessary to develop and sharply describe essential concepts which are called 'praxeological categories' or 'basic terms' which are interconnected into a deductive system. It should be possible then to deduce specific rules for specific endeavours from the general praxeological methodology (for example, of efficient organisation of people into teams). In this sense Praxeology is the scientific exploration of the conditions of efficient action and provides a typology of actions with respect to the concept of efficiency.

Logically the next step in the development of the theory is to establish certain relations between the once defined praxeological categories.

These relations are called praxeological principles. By far the most important is the "Wirtschaftlichkeitsprinzip" or "principle of economic efficiency"; that, given a certain expenditure of resources, the maximal realisation of the goal is attempted. According to Oskar Lange, the Wirtschaftlichkeitsprinzip can be used if the goal and the means of an action can at least be given on an ordinal scale. According to Oskar Lange, the Wirtschaftlichkeitsprinzip plays an important role in several disciplines, for example, in political economy, economy in general, organisation theory and in mathematical statistics, especially in the formulation given by Wald in 1950, where the Wirtschaftlichkeitsprinzip is defined as minimisation of expected loss.

Praxeology is primarily a normative approach. After observing and criticising the purposeful activities that are performed in the real world and trying to describe and classify them systematically, prescriptions or recommendations are produced the main aspect is not describing but changing reality.

Much of the Polish work is in praxeological organisation theory. Two books by Zielenieswki appeared in 1964 and 1969 directing attention to the success of an organisation as a whole and not of its members, which is similar to the approaches of Oskar Lange and Slutski in the field of political economy. Lange himself takes a socialist position; for him, a rational (he prefers "rational" to "efficient") method of production and distribution is only possible in a socialist economy, governed by a centrally managed plan. This, according to Lange, leads to a rationality of the economy as a whole which, in turn, leads to more efficient human behaviour in all fields of action. More recently (Gasparski and Szaniawski, 1976) Praxeology has been applied to the research process, to the analysis of goal determination in social development and to ecology which is similar to American writing in the framework of the so-called systems approach.

The methods used by Praxeologists are optimisation, mathematical programming, operations research and optimal decisions techniques. For

example, game theory is employed with reference to military considerations in the case of negative co-operation as it's called in Praxeology, or struggle. Some authors also use network and planning techniques, for example, PERT.

None of these techniques have been developed by Praxeologists. They are adapted from other disciplines.

My fourth task is to exhibit the core concept of Praxeology: efficiency, of which there are several definitions. The first is "effectiveness", which translates to "Wirksamkeit" in German and is synonymous with "goal orientation". Effectiveness is either graded or is zero-one-effectiveness (success or failure according to a predefined goal). The second definition is profitability, which is "Vorteilhaftigkeit" in German; (the difference relation of returns and costs of an action which is just taken). The third kind of efficiency is economic efficiency or "Wirtschaftlichkeit", which is defined as the ratio of cost and output. Besides these kinds of universal efficiency, Gasparski also considers efficiency in the physical, technological and economic sense, and attempts to show that physical efficiency is a special case of economic efficiency.

My fifth task is to relate Praxeology to other branches of science.

Praxeology is an interdisciplinary science, in common with General Systems Theory and Cybernetics (to a lesser extent Topology, Game Theory Measurement and Information Theory.

In its general orientation, Praxeology resembles General System Theory but emphasises its prescriptive element. Further, whilst the object of study in General System Theory is interpreted as a system, in Praxeology the object under study is either the act or the subject of action, called an agent.

From the point of view of Praxeology, Cybernetics is a supporting discipline (Hilfswissenschaft). According to Lange, Cybernetics serves primarily two purposes: 1. if human acts try to achieve a goal

indirectly, Cybernetics helps to analyse the long chain of causes and effects that were activated by human action; 2. if the conditions outside change during the act – especially if caused by this act, it is then necessary to change the means one employs to attain the desired goal. This gives rise to a new chain: Goal – means – change of conditions – new means – new changes of conditions, etc. Fast (and effective) adaptation to the new situation is only possible if accurate information on the changes of conditions is made rapidly available: (Lange, Politische Okonomie, Vol.I, p. 230f).

As already mentioned, Praxeology is pragmatic: it makes use of every technique that seems promising for achieving greater efficiency of action. Clearly, many techniques of formal, mathematical-statistical oriented decision theory are helpful.

But what is the status of behavioural decision theory? Opposed to decision theory in the former sense and opposed to Praxeology, which are both normative disciplines, behavioural decision theory is viewed as a descriptive approach, devoted to the question: How do humans solve problems/make decisions, and not: How should they solve problems/ decide among options in a rational/efficient way.

Here, again, Praxeology is flexible. If behavioural decision theory (as a part of psychology) finds out which kinds of errors and fallacies people usually commit, this is helpful for Praxeology in the sense that (efficient) directives can be deduced from such observations. In decision theory, we would call this de-biasing. A short discussion of this topic from a decision theory point of view is found in Moskowitz et al. (1976).

As pointed out by Rappoport and Wallsten (1972), the distinction between 'normative' and 'descriptive' approaches is not too clear and may obscure more than help. So, for example, so-called descriptive models may very well be used as a decision aid – the model of the judge instead of the judge himself. The bootstrapping approach – regression models of human inference behaviour – is a prototype of this (eg. Goldberg, 1968, 1970). Another example is the logical

decision tree approach of Kleinmuntz (1968) and others.

As far as I can see, these possibilities of aiding the decision maker
(and thereby leading to better decisions and more efficient actions)
have not been used or even mentioned by Praxeologists.

Are there fundamental differences in point of view between decision
theory, behavioural and statistical, on the one side, and Praxeology
on the other side? Obviously, Praxeology is concerned with goal-
directed acts as such, in a very wide sense, whereas in the decision
theory approach only a subclass of problems is considered. One such
restriction is directly to be taken from the term 'decision', which
means selection of a possible course of action out of a feasible set
of actions. In this sense, decision theory implies a rather high
structuring of the problem area under consideration. But many real-
life problems are not this well structured, or structurable. The
analysis of ill structured, nasty problems, has not been given this
much consideration. Therefore, Praxeology can claim a larger frame
of reference – at the expense of formalisations and quantitative
methods. The analysis carried out by Praxeologists is "conducted in
the traditional language of philosophy, sharpened by conventions
characteristic of logical positivism". But, in my view, not too
much! One example for this criticism: Lange (op.cit., p.241) speaks
of 'methods of historical and statistical verification' – which sounds
not too good in my ears. Certainly, there have been attempts to
formalise Praxeological concepts – the example of 'efficiency' has
been presented before – but the results are rather modest, as Gasparski
and Szaniawski (1976) confess themselves. They say: "However, they
(the attempts)   either consisted in simple adaptation of decision-
theoretic tools (...) or they fell short of their objective because of
the tremendous complexity of the concepts introduced."

Praxeology claims to constitute a 'logic' or 'grammar' of actions and,
in this sense, would be a well-defined theory, from which certain
theorems could be deduced (using standard logical tools). But when

we look at what Praxeology really is today, we see a rather verbal, philosophical essayistic body of definitions, classifications, declarations of intentions; a loosely organised frame-of-reference around the concept of 'efficiency of work, or acts'.

Only recently, Polish authors remember the great Polish tradition in logics and, consequently in view of the claims of Praxeology, try to give a logical account of Praxeology. This seems to me the more surprising, since in modern logics, there are most important developments which are not too far away from some Praxeological considerations - I have in mind the modern developments in epistemic and deontic logic. The second notion mentioned above, ie. 'grammar of actions', leads, of course, to theories of language: mathematical linguistics. Especially this aspect is taken up in the most interesting book by Maria Nowakowska (1973).

Nowakowska's approach is very well described in the preface to her (1973) book "Language of motivation and language of action": "The principal aim (...) was to construct a formal system in which behaviour is treated as a certain language, and analysed by methods of mathematical linguistics. (...) This idea was explored (...) by constructing a general system called the language of actions, with an embedded subsystem called the language of motivation.
In the first of these systems, actions (verbal or not) were identified with words of a certain vocabulary, while admissible strings of action (...) played the role of sentences in the language. In this way, the set of strings of actions (...) becomes formally identical with a certain language, and consequently, may be studied by appropriate methods of mathematical linguistics. (...) Outcomes of the strings of action may be treated in much the same way as the meaning of sentences, thus enriching the considered artificial language with a semantics."

Pask : Is it fairly true to say that Praxeology is the Polish word for Cybernetics, perhaps with a slightly left wing bias attached? Andre Couffignol defined Cybernetics as L'art d'assurer l'efficacité de l'action". It's only one definition (interdisciplinary fields are plagued by definition mongers) but it is a good one. Since it highlights a global approach to dealing with practical problems by looking

from the top to the bottom.

Schaefer : Perhaps Cybernetics is even more modest in its domain.

Pask : Cybernetics is probably more modest in what it _claims_ but not in its _domain_ of interpretation.

Schaefer : But Cybernetics has developed some techniques. I can't see any techniques developed by Praxeologists so it's a science just involving descriptions of terms like action, agent, and so forth, which are not more precise than you can find in a dictionary, and so what's the contents of this kind of science?

Since it's an old science, it started in 1913, I thought they must have found something, and I thought that the paper read at the Darmstadt Conference mentioned earlier just gave a very short overview of the topic and I was very sure that I could find real content in the articles and books cited but this was not the case.

SECTION 5

5. LOGICS OF ACTION FOR COMMAND AND CONTROL SYSTEMS

Overview

Throughout the Conference both decision task design and decision
taking are discussed as dynamic systems with strategic and tactical
components. It is thus appropriate to notice that a standard logic
with truth values "True/False" or Fuzzy or Modal truth values is not
able to accommodate process except by describing the process as
though it were some kind of pattern with an artificial expedient for
valuation (for example, giving the value "True" of "True/False" to
certain propositions or pattern descriptions which correspond to
actions that work or strategies that are successful). The artifice
is both arbitrary (how does one describe a process) and unnecessary,
since logics of action, pioneered, perhaps, by Von Wright who called
them Deontic logics, are quite well developed.

By the same token a decision system is, by its nature, bound to
receive and issue commands. Surely it is again possible to employ
expedients that image commands (or questions, for that matter) as
though they were the propositions of a propositional calculus or the
descriptions of a predicate calculus. But these devices are cumbersome,
inelegant, and, in the last resort, improper. Commands are not
propositional statements; moreover, they need not be treated as such;
for logics of command are also quite well developed (notably by Rescher).
Except for trivial cases (when, for example, the command is an
instruction necessarily obeyed by a computer), the naive, but prevalent,
assumption that command and control systems can be successfully analysed
and designed on the basis of inputs and outputs or the like is often
misleading. Such systems do not necessarily react like automata; if
they did, the concept of "decision" would reduce to no more than a
test leading to a response and the evidence contradicts this supposi-
tion: whatever else it may be, a decision is not a property test,
pure and simple, the outcome of which selects an action or option.
The several different conceptualisations of decision discussed at the
Conference all maintain that decision is far more than this.

It is, consequently, surprising that the analysis and design of systems for command and control (a fortiori, of decision) is usually stripped of the special logics intended to deal with these events.

The two papers in this Section, by Dr Kleiter and Dr Gelman, remedy this curious omission and provide two somewhat different (though entirely compatible) schemes that are the groundwork, at least, of the conceptual and mathematical framework able to accommodate the design and description of relevant man/man or man/machine decision systems. Dr Schaefer's paper might, incidentally, have been included in this Section insofar as recent developments in the "Praxeology" he describes take logics of action into account as a development peculiar to this science. The paper appears in the previous Section only because it is an historical survey of "Praxeology" which, Dr Schaefer maintains, has <u>not</u> until recently developed an identifiable theoretical structure of its own (and of which the notion of <u>process</u> is an essential part).

Dr Kleiter's paper contains a formal statement of a prototype action logic with some postulates derived from the psychology of decision making. It is precise but deliberately limited in scope. Dr Gelman's paper is phrased more generally since, in addition to his own ideas, the author is presenting the views and findings of a well established school of thought subscribing to a theory which has been applied quite widely. The two authors have different but (as before) compatible approaches, but both of them use the term "logic" seriously, to mean the <u>interpretation</u> of a language augmented by a calculus (rules for inference and derivation etc). The classical interpretation of a (propositional) calculus is a universal set, and the relation induced upon this set by interpreting a statement is a <u>model</u>; the logic is said to be "consistent" if all statements assigned the value <u>true</u> are interpretable as models in "all possible" universes (a thoroughly abstract idea), whereas "subsistence" is usually reserved for statement classes true in <u>some</u> universe or <u>some collection of universes</u> (though not necessarily in "all possible" universes of interpretation).

Other universes of interpretation exist (for example, the Herbrand Universes of classes employed for mechanical "theorem proving" in the predicate calculus). However, all of them may be fairly designated

174

"classical" insofar as they are static and the models obtained as
statement interpretations are "classical" insofar as they constitute
static relations. In contrast, logics of action are overtly or not,
"non classical" and their models are "non classical" because they
represent events or processes (arising from statements about
procedures which need not be serially ordered) and the universes of
interpretation constitute, in some sense, processors and possibly
different kinds of processors.

"Individual Goals, Plans, and Future Time Perspectives"

Kleiter : Over the last years we have observed an increasing interest
in ill defined and complex decision problems and also in the cognitive
processes underlying the behaviour while being engaged in such complex
problems. In multi stage decision making there is an exceptionally
interesting book by Bellman and his co-workers on psychiatric interviewing
where he treats ill defined problems by the use of adaptive dynamic
decision making models. But this book is, in part, speculative
and there are many gaps between this book and practical realisation.
On the other hand, the branch of psychology concerned with plans and
future time perspectives has been ignored by decision theorists
up to now and I thought it might be interesting for myself to have a
look at this branch of psychology.

The research done in future time perspectives has been restricted
exclusively to private affairs or personalised events, but I do not
think this is a necessary restriction.

My main impression, after quite an extensive examination of the literature,
is that there is a real need for a theory. Not only are the decision
problems ill-defined, but sometimes also the concepts used by research
workers in that field.

First, I shall give a brief overview of what has been done. In the
second part I'll try to give rather formal definitions of some of the
concepts which are involved in this research area and which I think may
be helpful in going on with the research or in connecting to decision
theory.

A final possibility is to improve some of the techniques used in
measuring plans and time horizons, cognitive representations and so on,
or to extract data and make predictions as in a computer guided
interview which is objective and is flexible enough to adjust to one
individual. Well, most of these applications belong to my own future

time perspective and nothing has been done about them.

NOTE : Dr Kleiter summarised and criticised existing work. In the
classical approach the dimensions of future time perspective (hence-
forward FTP) are shown in Fig 5.1. and typical characteristics of FTP
are indicated by the following quotations. "Future time perspective,
or FTP, is conceived as the timing and ordering of personalised future
events", or "future time perspectives might be regarded as that function
of the individual which permits him to draw up a model of the future."
The methods of research used at the moment are summarised in Fig 5.2.
and individual differences in FTP have been examined in terms of the
variables of Fig 5.3. Such investigations have met with varying success
but the findings are incoordinated and a cogent theory is lacking.

Kleiter : Let me come to the second part. Let's try to be a little bit
more specific about the language when we talk about such concepts. Let
me introduce a number of basic concepts (Fig 5.4.) To be complete P is
an individual but I usually omit P, since generally I'm talking about
one person. The state transition operator is U (Fig 5.4(3)) and the
operator I (of Fig 5.4(4)) was introduced by Von Wright. It is
important for the definition of an action. The essential distinction
is established in Definition 1 and Definition 2 of Fig 5.5. The
actual versus hypothetical state operator is used to define actions.
As the next thing we need time tags, and two types of time tags,
one an objective time tag, one at least subjective, and we need
objective time in respect to the subject himself. Well, I'll not
give many details about that, the system can be extended by being
more explicit about the metrics or about this time aspect, it can
be extended very directly. The next is we need some cognitive
representation of events of actions of the things we are going to
talk about, and, as a special kind of a cognitive representation,
an anticipation. For an anticipation we require only that it is
in, it is somewhere in the future. It's not necessary to have a
specific time tag, it's just if you introduce a specific point
now there is a direction after now.

The next two concepts are directly borrowed from decision theory.

Definition 3 (Fig 5.6) is "feasibility". We say that an action h
(which may be an action or a concatenated string of actions is
subjectively feasible if there is a cognitive representation, $R_{T}(h)$, of
that action. That is, R exists at an objective time (capital T) and
H is the action or action string. Refraining from action "Inaction"
is included as a possibility and we say an action is feasible if we
know there is a conditioning event, e, such that there is a probability
attached to the action given that conditioning event and this probability
is greater than zero. (Or you may introduce a flexible threshold of
"feasible up to a certain degree"). Similarly, if e belongs to a certain
class of events, then feasibility may be defined in respect to a certain
class of events.

In any plan there is at least one action which is feasible, otherwise
we wouldn't speak of a plan.

To summarise, feasibility may be characterised by the existence of the
conditioning event which gives the action string h a minimum chance of
being performed.

The next concept is expectation (Definition 4 of Fig. 5.6). Let there
be a cognitive representation RT(e) of an event (or action or action
strings). RT(e) is an expectation of an event e, at a certain time, T, if
the following conditions are fulfilled, first the type of the cognitive
representation isn't anticipation, it's somewhere in the future. Second,
there is a probability attached to that event. Finally, there may be a
specific time lag in the future, but it is not necessary that an
expectation is connected with specific time lag.

The next ingredient of the theory is "desire" (Definition 5 of Fig 5.7)
where utility enters the picture. If you expect something and you
connect with that some value we usually say this is a "desire", but you
may desire something (say that soccer team A wins against soccer team B)
and it has some utility or some value for you but you can't do anything
about it you are not a player.

Next consider the concept of a goal (Definition 6 of Fig 5.6). It is
difficult, as you can see, since we could say, "someone has got a goal".
Or "e is a goal". As it stands I use a deliberately weak definition to
avoid unwanted restrictions upon later development. This person
has a cognitive representation of the state of affairs he wants to attain,
and this e is a desire, and the probability of attaining that goal is
larger than a minimum. And last, but not least, if you feel that there
is something you can do, if you have some control. So "goal", as defined, is
a desire with some minimum chances of obtaining that goal and some
minimum degree of control.

There are different relations between goals (Definition 7 Fig 5.8).
A sub goal is a specific type of a goal defined in relation to a main
goal. G1 is a sub goal if it is a goal, if there is a second goal, the
main goal, and if the subjective time (t) of achieving C_T is before the
subjective time of achieving the main goal, and if the chances of achieving
the main goal are increased by having attained the sub goal beforehand.

One relation between goals is conflicting goals (Definition 8 of
Fig 5.8). Someone has conflicting goals if it's easier to attain a
first goal without the second and vice versa. Well, you can specify
that and translate it by probabilities. Let me give an example where
this possibility has sometimes been completely neglected; research
on delay of gratification. There has been a bit of discussion over
which concept is more general, "delay of gratification" or "future time
perspective, future oriented behaviour". It has been argued that all
kinds of future oriented behaviour involves delay of gratification
and usually the discussants neglect to note that "delay of gratification"
is a consequence only if there is also some timing conflict. Otherwise you
can usually attain a long term goal without being frustrated or without
delaying anything.

Another relationship between goals is that of compatibility (Definition 9
of Fig 5.9). Two goals may be completely incompatible (or incompatible
up to a certain degree).

The last relation (Definition 10 of Fig 5.9) is independence of goals.

In all these definitions everything is done by probabilities. Further
the system quickly gets terribly rich and complicated, though the
complexity is essential if the theory is to be realistically interpreted.

The next step is to use the definition of a goal to define a plan
(Definition 10 of Fig 5.10). We usually speak of a plan if we have,
at least, something feasible in our head where each of h is an
action string.

Pask : Sorry to interrupt, h is ordered in the subjective time tagged (t)
dimension, not necessarily in the action time tagged (T) dimension.

Kleiter : Yes, I should have used some kind of very weak order in
respect to priority.

Pask : Well, it need only be "prior to", in the sense that inside your
brain it will be a series of computations and you may have time tags
which are simultaneous. It's only when you get moving your arms or
shooting a gun or doing some other one of several mutually exclusive
and exhaustive actions that is to be tagged on to big T.

Kleiter : Yes, indeed, yes. But, as you can see, the definition of the
plan again is rather weak. It is not necessary that we have in our mind
a complete string of action. A plan may include a lot of gaps, for example.

Well, finally, regarding future time perspectives (Definition 12, shown
in Fig 5.11), it is proposed that we speak of future time perspectives
only if you introduce some time tags. If you have the minimum of timing.
The plan with a minimum of timing schedule for actions and goals. Any
questions? I have a reason why I am asking, because I have some
difficulties in going on. When writing this summary I hadn't
finished my paper completely. I intended to treat the question
of how we judge criteria, what good plans look like and other points
relevant to the Questionnaire. In the meantime I realised that these
issues are terribly difficult, especially for those tasks of interest.
We can speak of a good or bad plan according to the usual criteria of
consistency; a plan is good if someone plans, without logical

inconsistencies. From a substantive point of view a plan is not good
if it violates laws of nature or of anything else. But, in respect
to the future time perspectives, there seem to be problems peculiar to
decision making. If we consider tasks imposed by someone else, it is
not too difficult. But with personalised events, the tasks are self
imposed.

Zeidner : Why do you think a personalised task is self imposed?

Kleiter : The distinction is not clear and that is one of the problems
and it is a real problem in this framework.

Zeidner : You say the only kind of plan that you're interested in is
a plan with a minimum time schedule for actions and goals. Could you
elaborate on that concept?

Pask : I could possibly elucidate this a bit unless I'm interrupting
Kleiter in doing so. I think it's really a rather objective definition
of what you mean if you have a concurrent machine, such as brain (with
many t clocks, only locally synchronised) computing away, and you have a
man with a stop watch (a T clock only) outside, an external observer
looking at this thing. He is going to say that a goal is a class of
strings which constitutes a future time perspective. If he can put
(little)t's in synchronicity with certain of the (big) T ticks of the
stop watch and otherwise not. It doesn't matter how far away they are,
providing he can establish synchronisation. This, incidentally, is
another way of doing what Carl Adam Petri did some years ago and talked
about information transfer and concurrent nets, (establishing a
fundamental distinction between "control" and "information").

It is restrictive if you think of t as being one simple ordering.
But on the other hand, if you think of t (say the t_i s) as being the clocks
attached to a lot of loci of control in a net (I call it a command
graph) then there are definite moments when you can synchronise with the
internal process and when you can't: which allows both sharp valued
measurement and internal indeterminancy.

Broadbent : I think what's worrying me is, on the one extreme
one could have a plan which fits the definition of a plan but not the
definition of a future time perspective. For instance, where some
person has saved enough money to get married if and when he should meet
someone who wants to get married, but he has no prescribed time of this.

Kleiter : There's one possibility, for example, you can conceive
planning the activity of planning. For example, as an action at a
special point, where you plan to reconsider what you do.

Zeidner : That's where the synchronisation comes in.

Broadbent : What about the intervening case of a plan where there is a
series of actions one which must precede the other but no necessary
time. Isn't that quite interesting?

Pask : Consider the limiting case of two asynchronous creatures, rendered
locally synchronomous by an act of directing each other's attention, speaking
words which do the trick. Everybody would agree that we share some
programs, even if these are fuzzy programs. But a fuzzy program is
just an approximate way of talking about a fuzzy computation. In fact,
the proper way of talking about it, rather than by approximation, is
to say it represents a concurrent, not only parallel but concurrent,
process. A parallel process is one which isn't pre-empted by that scheme.
A concurrent one is one that might run into conflict or even into what
you call a conditioned inconsistency. Generalising still further,
using this analogy of dialogue that goes on inside, as well as outside
the brain, there's a good deal of evidence that one becomes aware of
things when there are bits of local synchronisity going on, and I would
submit, if that is so, we've now reached the limiting case where the
synchronisation is not between big T and little t, but between two of
the little t's, which I could label, for example, t_i .

Kleiter : We want to have a still closer look to the logic of
time, temporal logic, for example.

Pask : I agree, it's fascinating and absolutely essential. I came
up against the same problem myself over the question of behavioural
objectives and came to the conclusion that the only way to make sense
of the matter was by things called "command graphs". But "command
graphs" are disjuncts of Rescher's command statements and command
chains (such that no pre-empted chains are allowed). The only way we
found to deal with the business was to involve several of your
individuals "P" since commands are addressed to some "P". Given that,
it is possible to take advantage of Rescher's work. On distributive
commands, addressed to several people (or, perhaps, by generalisation)
to different parts of the same brain. That's perfectly reasonable,
you haven't given any particular specification of P. The Ps can either be
people in this room to whom I address a distributive command, "put
your hats on". Or it can be a number of loci of control in a computing
machine to which you might give a distributive command, "do so and so,
whatever is at your disposal, if you have that routine available,
execute it." This also allows for an extension of your compatibility
and your notion of independence and dependence.

Johnson : It's been a while since I read Toda, but isn't his cognitive
simulation, one way of looking at future time plans.

Vlek : Toda is rather vague. His language is sloppy compared to this,
but he is discussing all the possibilities that might occur in a complex
dynamic decision tree and how you deal with all these possibilities and
in fact he's also trying to name all the things that might occur to you
when you plan actions. I think that this might be very useful in
connection with that paper.

Pask : The key point is that systems that are a priori independent,
either because they are spatially distinct (in different processors) or
asynchronomous, can only be coupled together in certain ways to achieve
coherent action (depending on local synchronicity). The only way of
making a new system is by relating a priori independent systems. The
particulars of sharing together depend upon the kind of observation you
want to make; or whether you want to observe them at all, which may
not be a necessary requirement. You may only want to achieve agreement
(alias coherence) between the systems (for example, distinct decision
makers). Dr Kleiter's logic gives permissions for certain possible
and tractable modes of juxtaposition. It thus provides a formalism
in which to pose hypotheses in decisions and psychology in general.

DIMENSIONS OF FTP

1) Extension
2) Structure
3) Density
4) Directionality
5) Realism
6) Optimism
7) Certainty
8) Dominance of personalized events
9) Activity & Agency
10) Control
11) Location
12) Integration

Fig 5.1. Dimensions of FTP

1) Events
   self produced
   given events

2) Projective
   story completion
   TAT or TAT-Like

3) Questionnairs

4) Interview

5) Essay Writing

6) others like SD or graphical

Fig 5!2. Methods

185

1) Age
2) Sex
3) Social Class
4) Cross-cultural Differences
5) Delinquency
6) Psychotics & Neurotics
7) 'Closed Institutions'
8) Educational Influence
9) Social Relationships
10) Intelligence
11) Delay of Gratification
12) Achievement Motive
13) Actual Performance
14) Anxiety
15) Field Dependence
16) Activity
17) Optimism
18) Internal vs. external Control
19) Dogmatism
20) Test batteries & factor analysis

Fig 5.3. Individual Differences

1) P	INDIVIDUAL
2) $s_1, s_2, ...$	STATES
3) $s_1 \cup s_2$	TRANSITION OPERATOR
4) $s_1 I_p s_2$	ACTUAL / HYPOTHETICAL STATE IF P INTERFERES
5) $T(x)$	OBJECTIVE TIME TAG
6) $t_p(x)$	SUBJECTIVE TIME TAG
7) $R_p(x)$	COGNITIVE REPRESENTATION
$R_p^a(x)$	ANTICIPATION
8) $p(x)$	SUBJECTIVE PROB
9) $u(x)$	UTILITY
10) $c(x)$	DEGREE OF CONTROL

Fig 5.4 Basic Concepts

DEF. 1: EVENT

$$S_1 U S_2 \quad / S_1 \neq S_2 \text{ or } S_1 = S_2$$

DEF 2: ACTION

$$S_1 U (S_2 I S_3) , \quad S_2 \neq S_3$$

INPUT   ACTUAL OUTPUT   HYPOTH. OUTPUT

Fig 5.5 Basic Concepts

An action h is subjectively feasible iff:

1) there is a $R_T(h)$
2) there is a conditioning event e such that
    a) there is a $p(h|e)$ and
    b) $p(h|e) > f$, $f \geq 0$

> existence of a conditioning event which gives h a minimum chance of being performed

DEF 4: EXPECTATION

$R_T(e)$ is an expectation at T iff:

1) $R_T(e)$ is an $R_T^a(e)$
2) there is a $p_T(e)$
3) if there is a $t_T(e)$, then
      $t_T(e)$ is after $T(R^a(e))$

> anticipation of a more or less probable event with an optional time tag

Fig 5.6 Def. 3. Feasibility

$R_T(e)$ is a desire iff:

1) $R_T(e)$ is an expectation
2) there is a $u_T(e)$

---

expected event which is more or less valuable to P

DEF 6: GOAL

e is a goal of P at T iff:

1) there is a $R_T(e)$ and $R_T(e)$ is a desire
2) $p_T(e) > 0$
3) there is a $c_T(e)$ such that $c_T(e) > 0$

---

a desire with some minimum chances and some minimum degree of control

---

Fig 5.7 Def 5: Desire

DEF 7: SUBGOAL

G_1 is a subgoal iff:

1) G_1 is a goal
2) there is a goal G_2
3) there is a $t(G_1)$ and a $t(G_2)$
4) $t(G_1)$ is before $t(G_2)$
5) $p(G_2|G_1) > p(G_2|\text{non } G_1)$

is before a 'maingoal' and increases the chances to attain it

DEF 8: CONFLICTING GOALS

P has two conflicting goals iff:
there is a G_1 and a G_2 of P
such that

a) $p(G_1|\text{non } G_2) > p(G_1|G_2)$
and

b) $p(G_2|\text{non } G_1) > p(G_2|G_1)$

easier to attain G_1 without G_2 and
easier to attain G_2 without G_1

Fig 5.8   Def 7: Subgoal
Def 8: Conflicting Goals

DEF 9 : INCOMPATIBLE GOALS

G_1 and G_2 are incompatible for P iff:

$$1) \; p(G_1|G_2) = 0$$
$$2) \; p(G_2|G_1) = 0$$

DEF 10 : INDEPENDENT GOALS

G_1 and G_2 are independent iff:

$$1) \; p(G_1|G_2) = p(G_1|\text{non } G_2)$$
$$2) \; p(G_2|G_1) = p(G_2|\text{non } G_1)$$

further candidates :

time independence

value independence

independence in respect to control

independence of action strings

Fig 5.9. Def 9 Incompatible Goals

DEF 11: PLAN

$R_T(e)$ is a plan of P iff:

1) $R_T(e)$ is a $R_T^0(\bar{h})$

2) there is at least one action in $\bar{h}$ which is feasible

3) there is a goal of P such that
$$p(G|\bar{h}) > p(G|non\ \bar{h})$$

anticipated string of actions which increases the chances to attain a goal and with a minimum feasibility

DEF 12: FTP

$R_T(e)$ is a FTP of P iff:
1) $R_T(e)$ is a plan
2) there is a $t(h)$ and h is a member of $\bar{h}$ of that plan
3) there is a $t(G)$ and G is a goal of that plan

a plan with a minimum time schedule for actions & goals

Fig 5.10 Def 11; Plan, Def 12: FTP

"Some Methodological Aspects of Decision Making"

Gelman : My command of the English language is not too good;
consequently I have remained silent for most of the discussion and I
have listened carefully during the past few days and tried to gather
information and understand most of the reports.

On the first day I wasn't surprised to find that most of the reports
were devoted to psychological aspects of decision making. From my
point of view this is better than the mathematical theory of decision
making, because the modern formal logic on which the theory rests is
concerned mostly with the relations of symbolic forms and expressions
and can't analyse the structure of cognition and the process of thinking
as a whole.

Quite naturally, in view of experimental tradition, some psychologists
make short work of this problem in a different way, as, for example, in
using simple paradigms of choice and lotteries. The situation is bad insofar
as the areas of analysis of knowledge and cognition (or heuristics as
it is more modestly named) are very often treated as the sole concern
of the (traditional) psychologist.

This situation reminds me of the joke about somebody who has lost his
purse, and is searching for it under the nearest street lamp because
that is the only place that is lit.

The next day I was glad to hear about another street lamp - I mean
management games,etc ,complex decision tasks and cognitive processes,
and am convinced, once more, of the need for a proper special theory of
decision making. Of course, we must have such a theory if we don't
want to treat decision making as an intuitive skill; the times of
Machiavelli and Kissinger are ending and the US Army, as well as the
Soviet Army, must have at their disposal an effective theory of decision
making.

This theory will be devoted to representing decision making in its
own right, it will have its own special tools or methods, and what is

very important, will naturally unite the knowledge and methods of the psychological approach, the sociological approach, the economical approach, etc.

I cannot produce this theory from my pocket straight away. In the best traditions of modern science I can only sketch some possible ways of building the theory. But I _do_ claim that a general theory of activity is, in quite a natural way, a super (meta) system for embedding the theory of decision making. (As decision making is one of the crucial human actions in the system of human activity, Parkinson's Law operates, the first thing we must do is to find a larger system !!)

There is no need to give here a detailed review of the attempts to construct a theory of activity. Undoubtedly, the tradition goes back to the Greeks, to Aristotle. Many names should be mentioned, for example, paying tribute to M. Weber, who introduced the idea of the social action, and G.H. Mead, who used the notion of act of activity for treating cognitive psychological and sociological phenomena. Talcolt Parsons tried to unite these two approaches and constructed an analytical theory of social action in an endeavour to build a general theory of action. More recently, one should mention the Ford Foundation and its programme of developing behavioural sciences. But since these are well known, it seems more important for me to tell you briefly about one Soviet school which is connected with the name of Shchedrovitzky. He and his colleagues have been working in the field of the methodology of science since 1958. I will do my best to describe their approach to activity and its theory, at the same time using this as an opportunity to express my own point of view.

One of the most important prerequisites for the study of any objects and constructing their theory, are the methods and means of study used; firstly, the means for constructing the scope and subject matter of the theory with which we are singling out (and shaping) a definite fragment from the objective world for study and, next, the specification of its practical constructs.

I have referred to these practical constructs as gnoseological (epistemological) forms. In somewhat more restricted contexts they

are also known as Kantian a priori forms, or Kuhn's paradigms, or
Toulmin's ideals of natural order, or Bogdanov's organisational
forms, or Morganau's constructs, etc. Faults of many other approaches
to the study of activity arise from the lack of an effective
theoretical basis. Can the construct of activity be represented,
or constructed, in terms of the usual philosophical categories
like "thing", "attribute", or "process"? The answer is, no. Because
the activity isn't necessarily located in the man or near him. The
activity isn't a process. It is a system of processes, which form a
certain totality and could be decomposed into a certain set of parallel
processes.

The notion of process is shown in Fig 5.11. The process is a sequence
of any events that we can arrange consecutively in time. But activity
rarely is such a sequence. In general, activity is better approximated
by a pattern, as shown in Fig 5.12, because activity is a system of
parallel processes and the activity is the totality of these components.
But, if we take only separate parts, as shown in Fig 5.12, this does not
represent the notion of activity, only the whole system represents the
activity. In the structure which renders the activity a non trivial
system is a framework in which the past, the present, and the future,
are united.

The next point about the representation of activity is connected with
the fact that contrary to the usual approach, no attribute of a man or
a product of certain work can be taken by itself. Activity has no
personal character. It forms a certain universal wholeness which is
wider than the individuals who are the members of this system of
activity. An activity as a whole is not created by individuals. On
the contrary, it captures people. The German natural philosopher,
Humboldt, used to say that "it's not that the people are taking
possession of the language, on the contrary, language takes possession
of the people". (By the way, language is a special type of activity).
So, by being born a child enters a real existing world of activity,
and to become a man, he has to become "attached" to this system of
human activity, to learn certain types of activity, and, only after

that he becomes a person and a man.

In the present approach, activity is a real, existing, system with a rather heterogeneous set of components which include people, tools, material, machines, etc. All of these components are living and functioning according to their own laws. So to study the complex system (this system of systems or activity) we are in need of many task specific representations which depend upon the type of problem-solving tasks we are engaged in and many patterns for combining specific representations.

This is true, even in a simple case of a machine like a tape recorder we require and can use many types of representation, for example, a functional scheme, a technological scheme maybe a geometrical scheme or description, etc.

But in the case of activity (the system of action), the situation is more difficult and the main problem encountered in many attempts to construct a general theory of action, is connected with the lack of clearly stated means of encompassing special methods of representation.

This Russian school managed to develop useful means. Some of their notions are determined by the idea of reproduction of activity; others are connected with the specialised interpretations of activity in industry, in learning, in science, etc. There are many results but it is a pity that not all of them have been published. (Certain publication problems exist and the papers that are currently published are published in the Russian language, so I was asked when I was leaving Russia to help the school to make known their ideas in the West).

One of the representations of activity, a kind of functional representation, is sketched in Fig 5.13 which relates goals, tasks, means, methods and procedures with which the products are formed from initial materials. I confess that I am not prepared to give you more than its description as I received the invitation to this conference too late to prepare my paper properly, but that is my problem, and if there are any questions I shall do my best to answer them.

It is assumed that at the beginning of any activity, we have certain goals to fulfil. We need special skills for this activity. We really need special gnoseological (epistemological) forms to determine and formulate the tasks in which the skills operate. It seems to me that John Dewey, the American philosopher, was right when he said that we must be aware of the fact that problems are not given to the decision maker; he must take them up himself. So to extract the problems from reality, because reality is a mass of problems, a system of inter-connected problems according to Russell Ackoff, and to find out certain problems, you must have special means, ie. forms for organising forms.

What are they? I said before that there are certain gnoseological (epistemological) forms of organising individual or point-of-view specific or goal specific experience. Previously, I was thinking about the forms as a kind of a priori Kantian forms, but now in the light of modern contemporary science, it seems to me that the origin of these forms is ubiquitous and worthy of study. Only a part of them is connected to the structure of the brain. Other parts of the forms are connected with the development of man and the history of cognition. Sometimes it's the school and the environment you came from that determines them.

Modern science needs to make a special study of these forms if only because it can (very often) be shown that people use the wrong forms. Sometimes being in need of the new forms we have nothing to do but to use a mixture of the old ones, that is a question of skill. But if we want to do this on a scientific basis we must learn how to invent or to construct them.

Returning to the scheme of Fig 5.13. We have goal and task and, of course, you have seen there is a solid (non dotted) line from goal to task which implies amongst other things (dotted lines) that to determine the task we need and send for special means.

Similarly, to draw a solid line from an initial material to the product, we need a procedure that is based on a special method.

This scheme can be used for special discussions but unfortunately

time does not permit me to give you more examples of the fruitfulness
of this scheme. However, I am sure that such schemes are very
important because they provide new means for considering reality, and
can be used to review previous experience.

Another notion that has been used represents a "simple" act of activity
and gives us the possibility to decompose the whole system of activity
into such acts of activity.

As you can see, in the lower part of Fig 5.14 is represented the "objective"
part of activity: In - initial material that is transformed into
product P by a certain actions $G_i, ..., G_j$ done with the tools T (external
means) and knowledge as a kind of external means. On the upper part of
the scheme we have the "subject" part of activity: the actor, the
person performing the activity, the "screen" of his consciousness and
internal means.

Special place is taken by the goal of this act of activity since it can
be considered as both an objective or subjective element of activity.

Zeidner : How do you differentiate between a goal that exists
subjectively or objectively?

Gelman : It depends on a more general system in which this act of
activity is embedded. In this scheme (Fig 5.14) it is not important.
But when we are going to describe more complicated cases we have to
look where this goal is coming from. One basic tenet of System Research
is that we have to take into account at least two aspects of analysis;
that in one way our system is a component of a more general system and
in another way the system has its own components. From one point of view
the super-system dictates certain goals, and we have to find the place
of our system and the role it is playing in the general super-system.
But at the same time from another point of view the system can have its
own goals and can play its own games. It seems to me that Russell Ackoff
gives a fair account in his book "Planning of the Future" where a proper
accent is placed on the problems of humanisation, environment, etc.
that are involved. Let Fig 5.14 image a "simple" act (that is, simple from
some point of view). I am going to show how the notion of an act of

activity can be used for the representation more complicated types of activity with many actors.

Let's imagine that the goal of the act of activity (represented in the dotted area of Fig 5.15) consists in taking a decision. But the actor, a decision-maker, has no special procedures to fulfil the goal. So there is a need for another act of activity to elaborate suitable procedures. Of course, it is quite a different type of activity often done by another actor, a researcher (though it may be the same person, but if so, he is taking another position or point of view).

In the right corner of Fig 5.15 you can see the researcher analysing the existing experience (represented by several schemes of decision-making acts in the left-up corner of Fig 5.15 under the sign Σ). His goal is to supply a decision-maker with the necessary procedures and methods.

In this way it is possible to construct the representations of more and more complicated systems of co-operative activity and this has been done (the sketch of Fig 5.15 being refined for this purpose) to obtain interesting analytic results in many fields of cognitive and practical activity like Semiotics, Pedagogy, Industry, Architecture, etc. Some of the relevant original publications are :

1. Shchedrovitzky G.P. Methodological Problems of System Research. "General Systems", 1966, Vol XI, 27-53.
2. Shchedrovitzky G.P. Configuration as a Method of Structuring Complex Knowledge. "Systematics", 1971, Vol 8, No 4.
3. Shchedrovitzky G.P. Avtomatizatzia Proektirovania i Zadachi Razvitia Proektnoi Deiatelnosti. (Automatization of Design and Problems of Development of Design Activity - in Russian) Razrabotka i Vnedrenie Automatizirovanniikh Sistem v Proektirovanii. Moskva, Stroiizdat, 1975, 9-177.

Pask: In the hierarchy, which resembles a method we use for depicting what may be known in contrast to what may be done, we have found that recursive representations (under a "condense" operation) are also possible. Isn't that true for your colleagues and yourself?

Gelman : Maybe it will be useful to note here the basic difference
between gnoseological (epistemological) form as some means of con-
structing the subject matter of study and a model that substitutes
on the more higher level of research the initial subject of study.
It seems to me that the first problem is "shaping" a subject matter.
To do this, we need special gnoseological (epistemological) means.
The notions introduced earlier (of the act and system of activity)
appear to be very productive for constructing the subject matter of
the theory of activity.

Another problem arises later when studying the subject matter. If we
find out that for certain reasons we cannot deal with it directly,
then the subject matter is very often substituted by another object,
a model and the results of studying it are used for building up the
initial subject matter.

Pask : That is exactly the distinction between the entailment structure
(knowledge representation) and the "task structure" or "program graph"
interpreted as a (dynamic) modelling facility.

Gelman : But it seems quite clear that gnoseological (epistemological)
forms and models have different functions in cognition.

Pask : Obviously, I agree. The point is, however, that given a dynamic
theory of models with truth values related to action rather than (true,
false), the distinction, though needed in order to comprehend the
system, can be made at any convenient point. The definitions required
for making this distinction are recursive. Gergely tells me that an
adequate dynamic model theory has been developed by Montague. But
much the same story can be told in category theoretic terms, for
example, by Goguen, in terms of self referential logic (Varela, Maturana
and Von Foerster) or in terms of interpreted programming operations
(Gaines, Kaye, De Fanti, Winograd).

Vlek : This sounds like a development in problem solving. The question
is whether it is a more useful theory than the others.

Pask : Well, both are useful theories. The point is that theories
of action have the very odd property that they use non-classical
(ie. dynamic) models. The model is an activity not (as in the
classical case) a static relation induced upon a set by interpreting
an expression or executing a process.

Vlek : You know, Newell and Simon's point of view about the theory
of problem solving; a good theory enables us to write a simulation
program that imitates the behaviour. Or, even, a good theory is the
simulation program which contains a realistic mechanism which is
process varying for changing and adapting behaviour.

Pask : It's good if it's isomorphic up to simulation. But the
interpretations of such a theory are simulations, whereas what is being
talked about here is something isomorphic up to execution in reality,
rather than a simulation. For example, I can simulate a parallel
process in a machine but to actually do the parallel process I require
a different kind of computing device, a different domain of
interpretation, possibly several of them, to maintain independence
between several distinct processes. The distinction is probably
unimportant when there is no computational conflict. It is crucial
if computational conflict and parallelism are both permitted as, for
example, in an unrestricted Petri Net. If you like, such systems of
action involve concurrency not just (preordained) parallelism.
Exactly the same comments are valid at the level of cognition as well
as the level of action and of behaviour. Consider, for example, the
difference between simulating and actually doing the simultaneous
searches of associative memory (Norman and Lindsay's network or Quillian's
TLC, or Seeley Brown's procedural representations or the suite of
Kaye and Winograd programs forming KRL).

Vlek : We are talking about criteria for the usefulness of theories.
This one is useful because it refers to action rather than simulation.

Maybe we can pull things closer to home by talking about the
Bayesian model for information processing. Many people believed for

a long time that this might be a descriptive model if you change some
of its parameters or replace objective probabilities by subjective ones,
whereas later authors started to plea for developing a theory which
would meet the criterion of imitatability that is that it would contain
only concepts of operations which someone actually carrying out the task
could imitate in his behaviour. And that comes very close to simulation.

Pask : Close to simulation but it isn't quite the same. Simulation
is restricted insofar as it's got a serial order imposed upon it. This
would not matter too much if decision really _is_ made up of sequential
selections or if conditionalities _are_ reducible to tree structured choice
sets. I do not think this _is_ an adequate picture of decision but, even
if it _is_ adequate, it is still improper to make the a priori assumption
that this _is_ the case, to introduce it as an axiom of decisive systems
and, after that, to judge the adequacy of theories against the possibly
distorted (in my personal view very probably distorted) picture of
reality. By "distorted" I mean "unduly restricted", not utterly wrong-
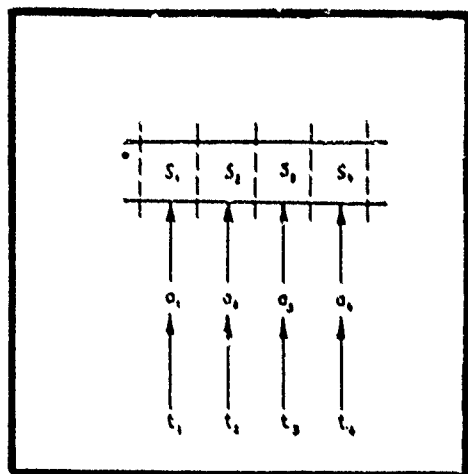headed or mistaken.

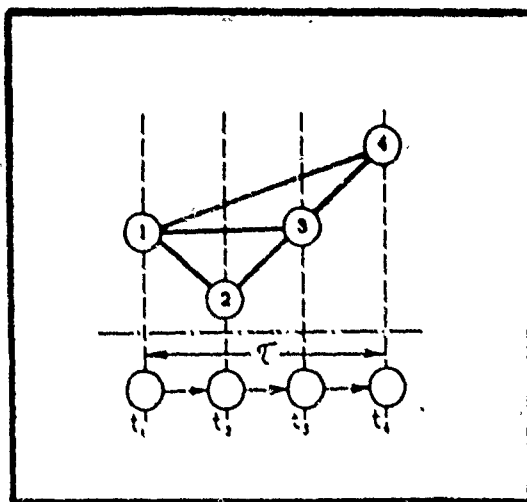Fig 5.12: (t = time instants, a = actions, s = states)

204

Fig 5.13: (t = multiply indexed time instants with
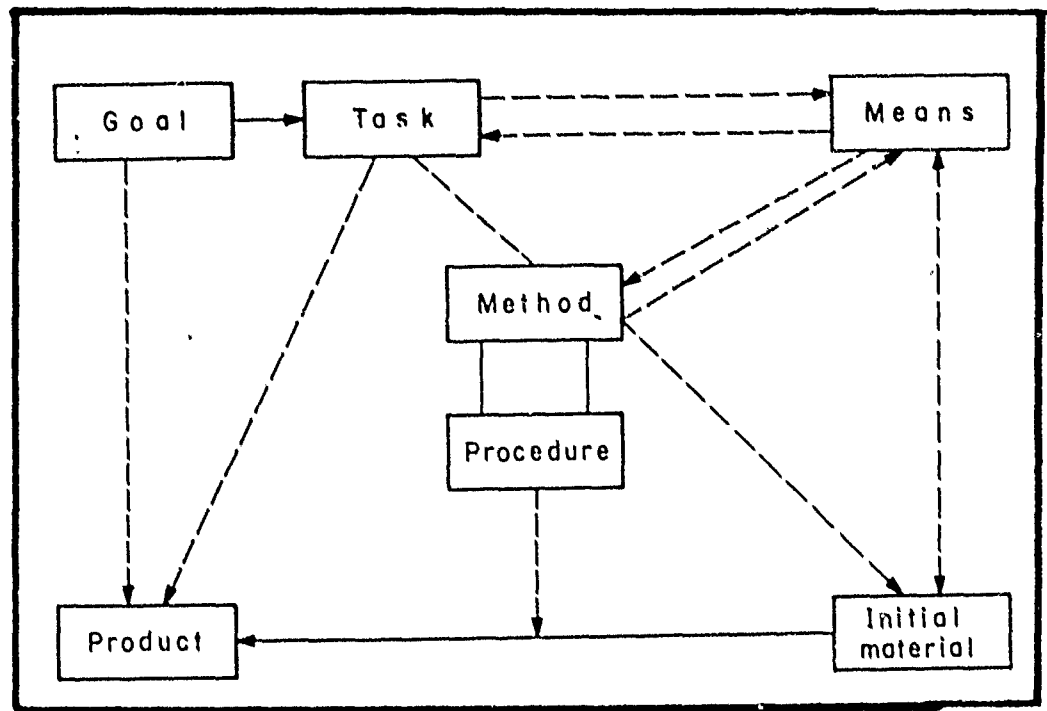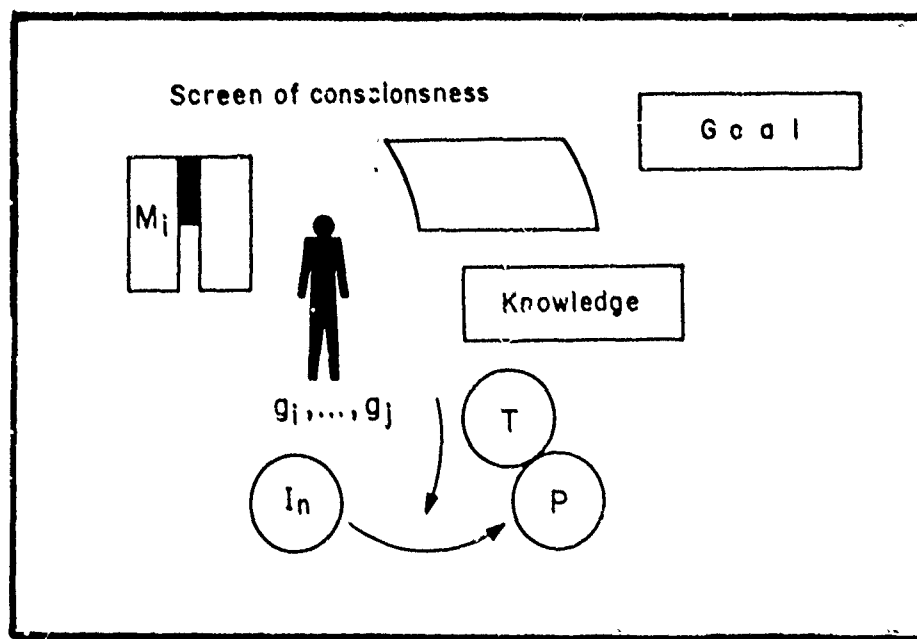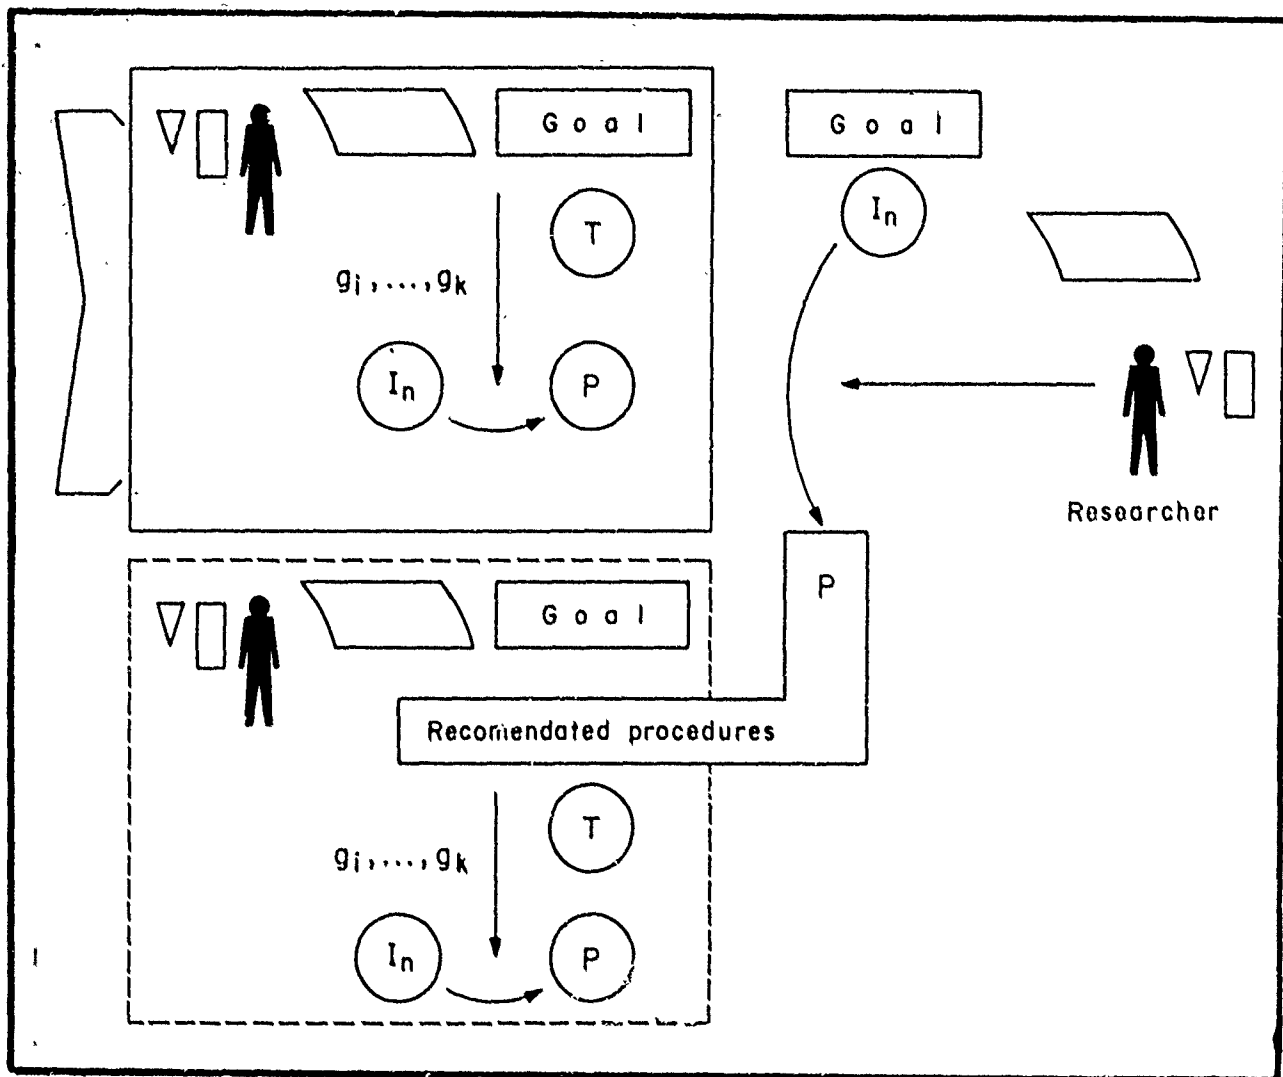τ = overall interval)

Fig 5.13

Fig 5.14

Fig 5.15

Section 6.

6. Conclusions reached and general recommendations.

Overview : On the evening before the last (morning) conference
session Dr Zeidner asked Dr Johnson and Dr Pask to summarise the
proceedings for presentation at the last session. (Summaries A
and B, below).

Regarding the original questions, the consensual view appeared to
be as follows :

   (a) Some individuals at least can be trained as decision makers,
but this should not be equated with training in probability assessment,
simple laboratory choice, judgement, or the like.

   (b) There is no unique answer to "how best to train them" but
quite a lot can be said about techniques that do not work on their own
(though they may contribute to a training process). The caveat of (a)
is important. Decision is a complex process, not a simple act as often
supposed in laboratory studies, so that training expedients like
providing immediate knowledge of results feedback relative to a particular
decision (for example, objective probability) do not teach decision even
though they improve certain relevant behaviours. Effective decision
training involves at least remedying mistakes and suggesting cogent
heuristics in terms of a common language established, perhaps, as a
compromise, and understood by the trainee and the instructor.

   (c) Decision making is probably a collection of skills integrated
in the framework of tasks between which there is adequate generalisation.

   (d) In specific situations decision skill (which, however
constituted, is complex and pervades the intellect) can be improved.
Probably there is sufficient task generalisation and transfer of learning
to comprehend classes of skills.

   (e) Decision makers must learn, preferably in context, concepts,
facts, principles, attitudes procedures and heuristics (all the
possibilities cited); some of them are usefully taught.

   (f) "Learning to Learn" applies to a decision situation (or class
of situations)."Learning to learn" involves (amongst other things)
formulating task representations and problems.

   (g) The "optimum balance" of aiding against training appears to
be task dependent and dependent upon idiosyncracies of decision style.

(h) Several coherent structuring schemes were discussed. A
further scheme is noted in Summary B. The environment may also be
personally structured by individuals who have 'learned to learn'.

It was felt that the "question set" should be enlarged to
include specific enquiries concerning

(i) Prescriptions for training that involve problem formulation
and recognition in large but specific task settings.

(j) Logics of action.

(k) Structuring, both internal and external.

Summary A .

Johnson : Learning is an invariable and unavoidable consequence of
decision making, probably of any other intellectual activity. If
individuals can be trained to be effective decision makers in both
unstructured and structured situations where they do learn by
experience, the problem is not whether they learn, but in fact how
to structure that learning and whether training can be explicitly
designed.

At least three kinds of notions are available. One is a traditional
approach, a rotation or structured sequence of job assignments. This
is an approach used currently in the military and in industry and in
some universities. An individual goes through a structured sequence
of job assignments, a career path, if you will, to arrive at a
certain position. Another approach is apprenticeship, much like the
graduate training that all of us went through. You study under one
master. He may not be able to tell you what he does but when you do
it like he does it then you've become a master yourself, (or a junior
master, depending on the circumstances). The third approach involves
explicit training design. If decision making training can be
explicitly designed, and I think that it can be, then how can that
training be accomplished most effectively. What are effective
instructional strategies for training decision makers. It is clear
that the state of the art is that such instructional strategies do
not currently exist. Thus the question is, how do you develop such
strategies.

The answer to the question depends upon the framework or perspective
from which we view the problem. For example, we need to consider
who is the decision maker, since it's not always clear. One of the
notions expressed was that decision making is a psychological waste
basket or, as Charles put it more smoothly, it's a pervasive concept
in psychology. I think that decision is a selection among alternatives
or among actions and implies the allocation of resources versus actions
which imply the commitment of resources. But I think we're concerned

with both aspects. The decision maker is an individual (or group
acting as a team) who controls or feels that he controls these
resources.

Let me sketch a framework for the problem of instructional
strategies. First there is a process which is very much like the
one that Charles (Vlek) described. I would use different steps,
but let me make some points about the kind of taxonomy that was
presented. A decision is really application of a process, a decision
process, by an individual or group of individuals to a task
environment, and this implies I think three classes of variables.
First , the process variables themselves, individual variables and
task variables. I think we're concerned with a subset of these ·
variables relevant to the case in which an individual is consciously
attempting to make a good decision. I don't think we're concerned
with such things as habits or those routine actions which are
carried out with very little conscious thought.

The first (process) variable indexes a continuum beginning with problem
recognition and ending at some other stage. It encompasses an arbitrary
number of sub-processes located at somewhat arbitrary points: the
particular set of tasks which were presented and discussed are also
somewhat arbitrary and the idea of occurrence in a sequence is also
somewhat arbitrary. There could be more or less task oriented skills
and whenever an individual finds himself performing in the role of a
decision maker, he may not explicitly run through this set of sub-
processes in a serial fashion, or may not even perform them all.

The process is directional, proceeding from problem recognition
through implementation of a decision or termination of the process
at some point, but also, as I think Gordon pointed out, an infinite
number of potential sequences of proceeding through this set of skills.
You may look, go back and obtain more information, reformulate your
problem, evaluate various consequences through some type of cognitive
modelling, then return to gather more information and there are
indefinitely many variables. Decision is also a purposive or goal

oriented process.

Last, but not least, our models are descriptive, in the sense that
we don't really know what a normative model of this process might be.
The descriptive model bears a curious similarity to an activity called
decision analysis which does purport to be a normative model. I think
it's clear from the description of the process that we think of
decision making as a collection of skills, although it's equally clear
that the composition of this collection of skills is somewhat arbitrary
and sketchy.

Variables of the second class are individual. There was relatively
little discussion of individual differences or of any moderating
factors of the decision maker himself. But I think it's clear that
there are large differences in decision style. For example, one
fashionable characterisation of "decision style" which we've used at
ARI in several studies, is in terms of three bi-polar dimensions:
"Abstract-concrete", which relates to the type of information a
decision maker prefers, "Logical-intuitive" which distinguishes the
form of information processing the decision maker goes through, and
"Active-passive" which describes the type of information acquisition
in terms of activity level. We found that these variables do, in fact,
relate very directly to how an individual attempts to solve a problem.
For example, if he's an abstract type, he may ignore the concrete
information and prefer abstract summary statements of information.
There is some experimental evidence that these are stable characteristics
which describe how individuals apply the skills of a decision process
to specific tasks.

The last class of variables, task description, is even less understood
than the individual variables. Only the Lens model, as mentioned by
Berndt (Brehmer), explicitly includes task variables or tast structure
as a formal component of decision making. A working hypothesis which
summarises several of the comments that have been made, is that the
decision process is applied to an internal representation of the task
environment, and that individuals may not be able to express their
knowledge of this representation of the task environment, although

they may be able to use their model in the task environment. The
distinction highlights the difference between a knowledge base
versus a performance base test of the task environment. This
point came up in discussion of the economics game (people playing
who couldn't express their knowledge of relationships, but could
apply their knowledge). One implication of this may be that the
task is represented in an analogue rather than a digital form, and
individuals may not be able to translate analogue models into a
serial verbal description.

Given this type of framework what are the implications for
instructional strategies? At first, if the decision process is,
in fact, a collection of skills, we can consider two types of cases.

One case is the training of individual skills; for example, the
calibration of probability assessment.

The second case is the training of a set or sub-set of these skills
as an integrated process, much as is done in most of the gaming research.

I assert, as another working hypothesis, that individual skills or
sub-processes, cannot be trained in the absence of task content. You
cannot calibrate somebody without referring to a specific type of
task. And one assumption that underlies this is that training should
focus on the individual applying this skill to a specific task. One
fact which led me to this type of hypothesis (that any type of feed-
back should be task related and not skill related), is the difference
between Larry's (Phillips) results in his calibration experiments
where the feedback was more skill related, ie. the feedback is based
on proper scoring rules. In contrast, for the studies that Dr Zeidner
mentioned, the feedback was task feedback. Subjects were not
calibrated and informed that their probability estimate is "off" by
so much from some norm, but rather the task itself was explained.

Another point relating to the learning of individual skills is the
learning of functional relationships within a task which may have
limited transferability to different functional forms. This is a
point which was made rather strongly, but I would put a different
interpretation on the results, ie. that subjects are learning a task

structure and not a decision process, the functional relations or relation between cues within a task. I think this has more implications for how a task structure is represented than it has for how people go about applying skills or making a decision for tasks in general. In other words, people may not be able to learn certain task relationships very well.

With respect to integration of skills, perhaps the only approach which has been tried, is something which might be called game base learning. However, aside from the fact that there is a lot of data indicating that the players in these games enjoy themselves, there is almost no learning data available which says they learn anything and further there are no associated instructional strategies to determine what they are meant to learn by playing a game. One possibility is to learn the relationships among the variables inherent in the specific system.

There is a representation of a task, much like Dr Kidd's economics model, but the generalisability of this learning would depend upon the fidelity of the game. To the extent that the game represents reality, the learning might be expected to transfer. To the extent that it is different, you might, in fact, have interference. With respect of the decision process within any game, there may exist different solution strategies, each strategy having a different degree of transferability to the more complex real world, so that generalisability would be dependent upon the strategy learned. A common expression is "players are playing the game rather than playing the problem itself, ie. the more general problem".

What they might learn, of course, is when to apply various skills and sub-processes of the decision process, somewhat like SR relationships or the TOTE concept of Miller, Gallanter and Pribram. This possibility is very important in the sense that it is common to find that people know a lot more than they apply. A striking example which occurs constantly in psychology is the use of statistics; people know much more statistics than they can meaningful apply. But they don't know when to apply what statistic, although if you lead them they do,

in fact, know the underlying statistic.

The question of the extent to which decision making skills can be
applied, in a larger or smaller context, depends I think upon two
factors. One is the process of decision or the skills themselves.
Next there is the representation of a task environment, which is
itself a skill. Now this latter factor implies that the traditional
concepts of transfer of training, for example, Thorndike's theory of
identical elements, may have some relevance. In any case, what is
it that determines the transfer of learning, from one context to
another.

A third important factor is, in fact, a process or skill of
articulating or of applying the decision process to a representation
of the environment. The question of applying decision skills in a
variety of contexts is akin to the concept of "learning sets", of
"learning to learn". One of the factors you "learn to learn" is how
to apply skills so that, for example, training in complex skills is
traditionally carried out by using a master/apprentice model. In part,
the absence of clear rules results in the fact that the skill must be
passed on by example, which may be why subjects can't verbalise,
because there are no clear rules to verbalise.

How then does learning take place? This relates back to the question
of a representation of a task, since a common explanation is that
sooner or later the novice acquires an internal representation of the
skill in question, and learns to recognise it when it appears. (It
"feels right" or it does not). It's much more an intuitive skill, a
"feel" rather than an explicit "yes, I know how to do this" or "no,
I do not" notion.

Hence, first of all, we might wish to teach decision makers how to
develop a description of the task environment (which, I think, would
contribute to generalisability), rather than teaching a description
of a specific task environment. Secondly, there is a decision process
(in a specific task environment usually called a procedure). I think

what we're attempting to teach is a process, how to go about applying the very flexible procedure which consists of a number of skills more or less of which may be applied in any given context.

In both cases, however, there is implied an interaction between the process of decision and some representation of the task environment.

Before considering what it is that the decision maker should be taught we need to delineate basic limitations of performance. For example, the number of alternatives a decision maker can simultaneously consider, which leads directly to the question of how should decision tasks be designed? It also leads to distinguishing deficiencies in performance which can, in fact, be overcome by training since it is not clear that all of the weakness in decision making can necessarily be remedied.

So what is it that individuals should or can be taught? Several points are relevant here. First, I think Ralf's (Schaefer) summary of Katrin's (Borcherding) talk implies that it's impossible to train decision makers to forecast random events, and that we need, if you will, a coherent task environment or a task representation before one can speak of training. Secondly, I don't think you can train either the specific skills or the process in general without some task content. We have to train people to do something. Thirdly, although contemporary instructional system design places great emphasis on analysing the job (what tasks have to be performed as part of the particular job, and what are the critical tasks or skills). So far as I know, very little work has been done with respect of decision making tasks. I don't think there is any clear idea of what it is that has to be done to make a decision.

There are some candidates, however. One is the notion of cognitive simulation or evaluating the consequences of alternatives. But now, how to train this? The only method which comes readily to mind is

that of exercise; you have people attempting to simulate the various consequences and pointing out the deficiencies, which is not a very powerful instructional strategy.

But the evaluation of the effectiveness of various forms of training is a difficult task itself if the performance of a decision maker is to be evaluated. What is a good decision even knowing the desired class of outcomes?

Regarding the state of the art, we have at least a kernel of a framework which we can use to develop some hypotheses about how to go about training, but I don't think we've uncovered any explicit framework for instructional strategies which can be directly applied to unsolved problems and I don't think we uncovered any problems which were solved. We don't have any basis for saying that we already know "crucial techniques" and are "going on to the next step".

The last point is recommendations for future research. I think two aspects come out as being very important. The first is problem recognition. You have to recognise a problem before you can begin to make a decision, or begin to apply a decision process. I don't think it's at all clear how we train or should train individuals to recognise that, in fact, a problem exists. Second, I think problem structure, while not mentioned very much, seems to be a critical element in the decision process, since if we refer back to some of the earlier work in problem solving, for example, Dunker's work Luchin's water jars once you've structured the problem appropriately the solution follows. Perhaps a critical element which is overlooked in most decision research is how problems should be structured. Most behavioural decision research presents subjects with a problem structure to begin with and looks at how that structure is articulated, not how that structure is arrived at. Part of problem structuring, and I think the task which is most amenable to research, but least understood, is how tasks are represented. What cognitive representation of tasks

do decision makers use or should they use. How should we train people
to articulate the elements of a task?

NOTE : Discussion by several members concerning the value, or lack
thereof, of short training courses on decision analysis.

Summary B (Revised)

Pask : "I address myself to training, learning and evaluation in
decision, of which evaluation is a slightly misbegotten word,
because it suggests we necessarily assign numbers to things.

My points put a spiral around what Ed (Johnson) was saying, though
'interestingly enough they were devised independently of Ed's comments
(this revision eliminates redundancy so far as possible).

The molar-molecular distinction in psychology and social science is
made between taking statistical measures, and developing mechanistic
or structural models, to represent what goes on. It's a rather sloppy
distinction and I believe that it has to be improved if we are to have
a theoretical basis in terms of which we can discuss the questions of
training and learning and posit decent hypotheses. (Incidentally, I
agree with Ed that learning is ubiquitous in all decision and that
hypotheses about training, and the appropriate formulation of 'decision'
includes the task environment, the actions to be taken and the cognitive
and conative processes, some of which are manifest as behaviour).

I would prefer to see a distinction of the kind in physics between a
macro-state description and a micro-state description. It should be
set up with the acknowledgement that the phenomena under scrutiny are
very much more complex than those encountered in the physical sciences
and also to countenance phenomena of a different type. In decision we
are interested sometimes in the measurements of an outside observer
(sharp valued micro-statements about cogent choice; statistical state-
ments about values of macro-variables). We are interested always in
the statements and conditions of participants in the decision, who reach
agreements on the level of micro-level considerations (of methods,
heuristics and the like) and entertain personal macro-state descriptions
(of their belief, doubt and the like, expressed numerically, or verbally
or in other ways). Some speakers have made these distinctions (micro/
macro and participant/observer) quite explicitly but there is not, at
this stage, an adequate framework in which to give precise expression
to the superficially self-evident but truly very elusive concepts of
decision, let alone to manipulate and control the decision process.

For example, how do you make sense of the idea that decision is a process involving a task structure, a decision making individual or group and one or more internal representations of the task (or tasks) in hand?

So far as I can see, something of this kind is required.

The environment, actually dynamic, can be given a static representation (in fact, an indefinite number of them) in which tasks, themselves structured, will be related and in which the decision agents will be represented as though they were static. As seen by an external observer, or by a participant who momentarily adopts the stance of an observer, the representation has the micro structure of a cyclic mesh. Its macro structure entails many varieties of doubt on the part of an observer; for example, doubt regarding the perspective to adopt in dealing with the environment (or guessing what perspective someone will adopt), doubt regarding which submesh of tasks is relevant and doubt, given relevance (or problem formulation) about how to solve a problem. Insofar as such meshes are used to depict 'internal descriptions' it may be desirable and possible to constrain them so that their structure satisfies norms: constraints like this could be imposed, perhaps, by training which is task specific but also integrative, leading to generalisation.

However, the static picture is a common (possibly an unavoidable) fiction. In fact, the agents and their environment are dynamic entities. Seen impartially, whatever that means, the mesh is (continually) unfolding into procedures for thought and action as directional, and in a limiting case tree like, structures oriented with respect to one or more perspective.

These unfoldings are decision in the broadest sense, viewed microscopically.
The diversity and possible incoherency of these unfoldings gives rise to doubt indexed by macroscopic variables.

Under these circumstances, it becomes essential to consider different varieties of doubt on the agent's part: for example, doubt about a

possibly deterministic past (Broadbent's memory related uncertainties) and the future of plans and prospects (Kleiter).

The participating agent always sees things like this and it seems that, as a conference, we countenance the following possibilities. If the deciding agent is an individual, he may have one or many perspectives, though frequently one focus of attention is dominant. If the deciding agent is a group then many perspectives exist but may be unified by agreement. In either case, the static mesh is simply a way of speaking (fictionally, but conveniently) of classes of iterative unfoldings (an epistemological domain which _is_ the 'context' called for as the prerequisite of decision). So far as the deciding agent is concerned, all decision is dynamic; the future and past uncertainties are subjective. The debate in Section 5 bears upon how often an individual agent maintains more than one perspective at once.

The framework I have suggested does remove the paradoxical overtones of "who is the decision maker" and it does help in the debate about whether to train people or to provide idiosyncratic decision aids to suit their mental quirks and oddities. It also provides a more practicable foundation for decisions in really complex systems than the elegant, but locally and specifically applicable, Lens Model. As a matter of fact, many of you seemed to have such a framework in mind though, with few exceptions, you were diffident about discussing it.

The diffidence is quite reasonable. This picture is typical of Cybernetics and General System Theory (possibly, later developments of Schaefer's Praxeology). Until recently, the logical and mathematical tools for handling this type of framework have not been available and the empirical investigation it engenders (for example, exteriorising lengthy stretches of cognition as many typed behaviours) are quite alien to the highly controlled, short session, laboratory experimental paradigms. I do not think these tools are yet fully assembled but enough of them have been manufactured and used successfully to form a competent tool kit. The actual state of the art in respect to the mathematical and empirical tools of Cybernetics or General System Theory is probably in advance of the general estimate. This comment is based upon information you are unlikely to take account of unless you are immersed in this field, and caution, but no longer diffidence, is needed in applying these methods. My claim is dogmatic because I

am more deeply immersed in these fields than many of you and because
of the obvious relevance of these methods to the matters discussed.

For example, the unfolding aspect,'decision' itself, is open to
scrutiny in terms of the action logics of Kleiter and Gelman: others
of like kind exist, with the non classical interpretations which
prove essential in the evaluation of action. Vlek's cognitive maps
(and most people's task structures, notably Kidd's and Johnson's) are
identifiable as the procedural meshes that represent dynamism as
though it were statically inscribed. Novel concepts of stability
such as "Organisational Closure" overcome many of the difficulties in
identifying those systems that do and those that do not have a chance
of working: it seems likely that Johnson's stylistic categories and
Broadbent's memory types could be more readily related to complex
processes if expressed in these terms. Various extrapolations of
"coherence truth" provide an initial calculus for agreement and there
are several candidate unification schemes. (In this paragraph I refer
to the work of Von Foerster, Varela, Goguen, Gaines, Gergely and Nemeti).

All this suggests, of course, that the list of questions (a to h in
the Introduction) be extended to cover such issues as 'logics of
action' and the usefulness of "category theoretic representation".

My last point relates individual, stylistic and cultural differences
to one of the recommended training expedients "learning to learn".
Whilst generally agreed to be a good thing, how do you 'teach people
to learn' and what does the art of 'learning to learn' comprise.

Even in the absence of a systemic picture, someone who 'learns to
learn' must recognise their personal proclivites and competences and
be able to adapt them, so far as can be done, to the task in hand
(whether these are individual memory expedients or individual choice
of heuristics). The point is clarified by phrasing the matter
systemically, since, given a representation of what may be known and
done, it is possible to distinguish between the strategies adopted
from competence in using them: thus, for example, to recommend

training strategies matched to the stylistic competence of the trainees.

But more is gained by a systemic view. Over and above the acquisition and use of self knowledge, at least two other component skills are needed. In the first place, someone who has learned to learn is able to structure an otherwise unordered or hopelessly complex situation so that he can learn about it. This skill is trainable at any rate for the majority of subjects, and one important facet of the skill (upon which positive transfer also depends) is ability to use and not misuse analogy relations that are given as valid, or to use original analogical reasoning as an innovative, (some people called it "creative") backbone for decision.

The other skill, the less tractable of the two, but, all the same, indoctrinable, is simple "responsibility" especially for the structure an individual has imposed upon the environment or situation. Persuasion techniques are quite helpful in establishing the contract of responsibility but, like Dr Zeidner, I should be suspicious of them unless their mechanism could be incorporated as part of the general system model. Fortunately, there _is_ some mechanism behind the method. "Man" (in roughly the sense of Kelly's "Man as an experimenter") appears in the systemic scheme as a coherent system of beliefs, alias, a coherent set of stable procedures undergoing execution and development (_this_ is _not_ the only definition of _man_ of course, just a definition in the common language of a theory).

From listening, it seems to me that you all feel the need, as Dr Zeidner does, for a theory of decision applicable to complex systems. Though the terminology has varied throughout I am inclined to believe we are nearer than might be supposed to theories of the kind that are required.

Note : Since the session was over-running and several participants had to leave, there was little formal discussion of these summaries.